

Efficient Distributed Online Prediction and Stochastic Optimization with Approximate Distributed Averaging

Konstantinos I. Tsianos and Michael G. Rabbat

Department of Electrical and Computer Engineering

McGill University, Montréal, Québec, Canada

Email: konstantinos.tsianos@gmail.com, michael.rabbat@mcgill.ca

March 6, 2014

Abstract

We study distributed methods for online prediction and stochastic optimization. Our approach is iterative: in each round nodes first perform local computations and then communicate in order to aggregate information and synchronize their decision variables. Synchronization is accomplished through the use of a distributed averaging protocol. When an exact distributed averaging protocol is used, it is known that the optimal regret bound of $\mathcal{O}(\sqrt{m})$ can be achieved using the distributed mini-batch algorithm of Dekel et al. (2012), where m is the total number of samples processed across the network. We focus on methods using approximate distributed averaging protocols and show that the optimal regret bound can also be achieved in this setting. In particular, we propose a gossip-based optimization method which achieves the optimal regret bound. The amount of communication required depends on the network topology through the second largest eigenvalue of the transition matrix of a random walk on the network. In the setting of stochastic optimization, the proposed gossip-based approach achieves nearly-linear scaling: the optimization error is guaranteed to be no more than ϵ after $\mathcal{O}(\frac{1}{n\epsilon^2})$ rounds, each of which involves $\mathcal{O}(\log n)$ gossip iterations, when nodes communicate over a well-connected graph. This scaling law is also observed in numerical experiments on a cluster.

1 Introduction

In order to scale up to very large optimization problems, it is necessary to use distributed methods. Such large-scale problems frequently arise in the context of machine learning. For example, when one wishes to fit a model to a very large (static) dataset, the data can be partitioned across multiple machines and the aim is to obtain an accurate solution as fast as possible by leveraging parallel processing. Alternatively, when data samples arrive to the system in a streaming manner, as in the setting of online learning, and the model is progressively updated after each sample is observed, a distributed system may be required when the rate at which samples arrive is faster than a single machine can handle.

In either the static or online setting, the nodes in the distributed system must communicate in order to coordinate their computation, and it is important to understand the extent to which this communication affects the overall scaling performance of any method. In many algorithms, each node maintains and updates its own copy of the decision variables, and communication is performed periodically to synchronize the variables across the network. Typically this synchronization operation simplifies to distributed averaging: each node initially holds a local vector and the goal is for all nodes to compute (exactly or approximately) the average of these vectors.

This paper proposes and studies a method for distributed online prediction and stochastic optimization using approximate distributed averaging to coordinate and synchronize the values at different nodes. Previous work has shown that methods using exact distributed averaging achieve the optimal regret bounds for distributed online prediction. When an approximate distributed averaging protocol is used, there is some residual synchronization error in each node's decision variables. The magnitude of this error is inversely

related to the amount of communication. We present a general framework for online prediction and stochastic optimization using approximate distributed averaging protocols, and we characterize the performance requirements (in terms of accuracy and latency) that a distributed averaging protocol must satisfy in order to achieve optimal rates. Then we focus on gossip protocols for distributed averaging and precisely characterize the number of gossip iterations that must be performed in order to control the error and meet the requirements mentioned above. Gossip protocols have a number of attractive features, including that they are simple to implement, robust to communication failures, and they can naturally be implemented in an asynchronous manner. For this reason, gossip-based distributed optimization methods have recently received considerable attention [5, 10, 13, 17, 24, 33].

Before giving a detailed statement of our contributions, we describe the problems of online prediction and stochastic optimization.

1.1 Online Prediction

In online prediction, a system must sequentially make predictions. We consider the setting where the set of possible predictors \mathcal{W} is a compact convex subset of a real Euclidean vector space. After the system makes a prediction $w(t) \in \mathcal{W}$ for $t \geq 1$, a sample $x(t) \in \mathcal{X}$ is revealed, and the system suffers a loss $f(w(t), x(t))$. The random samples $x(t)$ are assumed to be drawn independently and identically from an unknown distribution over \mathcal{X} . After observing $x(t)$ and suffering the corresponding loss, the system updates its predictor to $w(t+1)$, receives a new sample $x(t+1)$, suffers the additional loss $f(w(t+1), x(t+1))$, and the process continues.

The goal of an online prediction system is to minimize the loss accumulated over time. Specifically, after a total of m samples have been processed, the regret accrued by a single-processor online prediction system is given by

$$R_1(m) \stackrel{\text{def}}{=} \sum_{t=1}^m [f(w(t), x(t)) - f(w^*, x(t))] , \quad (1)$$

where $w^* \stackrel{\text{def}}{=} \arg \min_{w \in \mathcal{W}} \mathbb{E}_x[f(w, x)]$ is the predictor that minimizes the expected regret with respect to the unknown distribution over the samples $\{x(t)\}_{t \geq 1}$. When the loss function is convex and smooth (i.e., f has Lipschitz-continuous gradients), it is known that no first-order method can achieve a bound on the expected regret which is better than $\mathbb{E}[R_1(m)] = \mathcal{O}(\sqrt{m})$ [4, 18, 28]. At the same time, this performance is achieved by a number of methods [28], including Nesterov's dual averaging algorithm [20].

Applications of online prediction methods are typically those where predictions must be made in real-time when a query arrives to the system. When queries arrive at a rate which is faster than a single processor can handle then a distributed online prediction system must be used. Consider a general system of n nodes. Each node receives a proportion (say $1/n$) of the queries arriving to the system and makes predictions. After making a prediction $w_i(t)$, node i observes a sample $x_i(t)$ and suffers a loss $f(w_i(t), x_i(t))$. The nodes can individually update their predictors based on their local observations (e.g., using $\nabla_w f(w_i(t), x_i(t))$), and they can also communicate with other nodes to obtain information about their predictions and observations.

Consider a general framework where, in the t -th round, the network collectively makes a batch of $b \geq n$ predictions $\{w_i(t, s) : 1 \leq i \leq n, 1 \leq s \leq b\}$, where b is an integer multiple of n . After making the prediction $w_i(t, s)$, node i receives an observation $x_i(t, s)$, and suffers the corresponding loss. Once the entire batch is processed, the nodes perform some communication, and then the next round begins. During the time that the nodes are communicating, additional queries may arrive to the network, and so the nodes continue to make predictions using their most recent predictor. Let μ denote the number of additional samples that arrive to the system while nodes are communicating (μ/n samples arrive at each node). For the general framework just described, performance is measured in terms of the regret

$$R_n(m) \stackrel{\text{def}}{=} \sum_{i=1}^n \sum_{t=1}^{\frac{m}{b+\mu}} \left(\sum_{s=1}^{\frac{b}{n}} [f(w_i(t, s), x_i(t, s)) - f(w^*, x_i(t, s))] + \sum_{s'=1}^{\frac{\mu}{n}} [f(w_i(t, s), x_i(t, s')) - f(w^*, x_i(t, s'))] \right) . \quad (2)$$

The terms in the sum over s correspond to the regret suffered while locally processing samples which will be used to update the prediction, and we have precise control over the number of terms in this sum since we can specify b . On the other hand, the terms in the sum over s' correspond to the regret accumulated while

communication takes place, and the only way to control this portion of the regret is to limit the time spent communicating.

Clearly, we cannot hope to obtain a better regret bound using n processors than would be obtained using a single processor; i.e., $\mathbb{E}[R_n(m)] = \mathcal{O}(\sqrt{m})$.¹ At the other extreme, consider a naïve approach where nodes do not communicate at all. In this case, $\mu = 0$. However, since each node operates in isolation, it individually accrues a regret which is bounded by $R_1(m/n)$, the regret that a single processor would have accrued using only m/n samples. Thus, the no-communication approach has an overall regret bound of $\mathbb{E}[R_n(m)] = \mathcal{O}(n\sqrt{m/n}) = \mathcal{O}(\sqrt{nm})$, which is a factor of \sqrt{n} worse than what we hope to achieve [8].

It has been shown that the constant hidden in the $\mathcal{O}(\sqrt{m})$ bounds for serial online prediction algorithms is proportional to the standard deviation of the random gradients $\nabla_w f(w, x)$ [16, 35]. Based on this observation, Dekel et al. [8] propose the distributed mini-batch algorithm. Specifically, the predictor at each node is held fixed at $w_i(t, 1)$ for all predictions in the t th round. Then, the gradients $\{\nabla_w f(w_i(t, 1), x_i(t, s)) : 1 \leq i \leq n, 1 \leq s \leq b/n\}$ can be averaged to reduce the standard deviation. Averaging the mini-batch of b gradients across all nodes involves communication, and Dekel et al. [8] propose to use an exact distributed averaging protocol. Once the protocol returns, all nodes hold the exact average gradient,

$$\frac{1}{b} \sum_{i=1}^n \sum_{s=1}^{b/n} \nabla_w f(w_i(t, 1), x_i(t, s)) .$$

This distributed mini-batch procedure with exact distributed averaging is shown in [8] to achieve a regret bound of

$$\mathbb{E}[R_n(m)] = \mathcal{O} \left(b + \mu + \sqrt{\frac{b + \mu}{b} \cdot m} \right) .$$

Thus, by choosing b appropriately (i.e., taking $b = \Theta(m^\rho)$ for any $\rho \in (0, 1/2)$), the distributed mini-batch algorithm achieves an expected regret of $\mathbb{E}[R_n(m)] = \mathcal{O}(\sqrt{m})$.

Although exact distributed averaging enables this approach to achieve the optimal regret rate, it also has some drawbacks. In particular, exact distributed averaging generally requires blocking communications or some form of barrier synchronization across the network. All nodes remain blocked until the operation completes, and consequently, the computation proceeds at the pace of the slowest node. This is undesirable in clusters, where infrastructure is prone to fail and where resources are typically shared by many users. Motivated by this observation, the aim of the present manuscript is to develop distributed methods which use approximate distributed averaging protocols that can be much more flexible and can be implemented to run asynchronously.

1.2 Stochastic Optimization

The problem of stochastic optimization is closely related to that of online prediction. In stochastic optimization we seek to find a value $w \in \mathcal{W}$ that solves the problem

$$\begin{aligned} & \text{minimize} && F(w) \stackrel{\text{def}}{=} \mathbb{E}_x[f(w, x)] \\ & \text{subject to} && w \in \mathcal{W} . \end{aligned}$$

A serial first-order stochastic optimization method sequentially makes stochastic gradient evaluations, observing $\nabla_w f(w(t), x(t))$, and uses these to update the decision variable $w(t)$. The performance of a single-processor method is measured by the optimality gap $\Delta_1(T) = \mathbb{E}[F(w(T))] - F(w^*)$ after T updates.² When $f(w, x)$ is convex in w for all $x \in \mathcal{X}$ and when the $x(t)$ are i.i.d., it is known [35] that any online prediction algorithm that has an expected regret bound of $\mathbb{E}[R_1(m)] = \mathcal{O}(\sqrt{m})$ can be used for stochastic optimization and will achieve an optimality gap of $\Delta_1(T) \leq \frac{1}{T} \mathbb{E}[R_1(T)] = \mathcal{O}\left(\frac{1}{\sqrt{T}}\right)$.

¹One might hope to achieve this rate if communication was arbitrarily faster than the rate at which queries arrive to the network and are processed, in which case every processor could always share every sample with every other processor and they behave overall as a single processor observing all m samples.

²Note that the expectation is taken in the definition of optimality gap because the predictor $w(T)$ is random since it is a function of the random observations $x(t)$, $t = 1, \dots, T$.

Stochastic optimization is useful when one may wish to fit a model w to a very large collection of data $\{x_i\}_{i=1}^N$ by minimizing the average loss $\frac{1}{N} \sum_{i=1}^N f(w, x_i)$ over the entire data set. Rather than performing gradient updates using all of the data, which can be very time-consuming when N is large, a stochastic gradient is obtained by drawing one sample $x(t) = x_i$ independently and uniformly from the collection at each iteration, and using $\nabla_w f(w(t), x(t))$ as a surrogate for $\frac{1}{N} \sum_{i=1}^N \nabla_w f(w(t), x_i)$. Clearly, when $x(t)$ is drawn independently and uniformly over $\{x_i\}_{i=1}^N$, we have that

$$\mathbb{E}_{x(t)}[f(w(t), x(t))] = \frac{1}{N} \sum_{i=1}^N \nabla_w f(w(t), x_i) .$$

The motivation to use a distributed system for solving stochastic optimization problems is often to obtain a solution faster than a single processor. In contrast to the online prediction setup, the system does not accrue any additional regret (or suffer an increase in the optimality gap) when communicating. On the other hand, it is of interest to study the scaling properties of distributed methods. For an optimal single-processor method to ensure the optimality gap is bounded by $\Delta_1(T) \leq \epsilon$ for some $\epsilon > 0$, it must evaluate $T \propto \frac{1}{\epsilon^2}$ stochastic gradients. An ideal distributed method using n nodes would obtain a linear speedup by guaranteeing the same level of accuracy after a time equivalent to what it would take the single processor to evaluate $\frac{1}{n\epsilon^2}$ stochastic gradients. To achieve this runtime scaling, a distributed method must not require excessive communication per round, relative to the time spent computing.

1.3 Contributions

This paper proposes and analyzes a method for online prediction and stochastic optimization. The proposed method builds on the dual averaging algorithm and the idea of using distributed mini-batch computations. Specifically, we consider a framework using general approximate distributed averaging protocols that ensure the dual variables at each node are synchronized to within accuracy δ after a latency $\mu(\delta)$. Our first contribution is a regret bound for this distributed framework which enables us to determine the relationship between δ , $\mu(\delta)$, and the mini-batch size b that is sufficient to ensure the optimal $\mathcal{O}(\sqrt{m})$ asymptotic regret bound is achieved.

Then, we focus specifically on using gossip protocols for approximate distributed averaging [3, 9]. Our second contribution is a bound on the number of gossip iterations required to satisfy the relationship between the accuracy δ , latency $\mu(\delta)$, and mini-batch size b , so that the optimal regret bound is achieved. In particular, the number of iterations scales as $\mathcal{O}\left(\frac{\log(n)}{1-\lambda_2(P)}\right)$, where $\lambda_2(P)$ is the second largest eigenvalue of the transition matrix P of a random walk on the network topology over which messages are passed. When the network is well connected (e.g., when it is an expander [25]), then $\frac{1}{1-\lambda_2(P)} = \Theta(1)$ as $n \rightarrow \infty$, and so the number of gossip iterations (which is proportional to the time spent communicating) is $\mathcal{O}(\log(n))$, which is no more than the time spent communicating per round if an exact distributed averaging protocol were used, up to constant factors.

Finally, we study the proposed approach in the setting of stochastic optimization. Specifically, we derive bounds on the optimality gap as a function of the accuracy δ . We show that the number of rounds required to drive the optimality gap below ϵ (i.e., $\Delta_n(T) \leq \epsilon$) is of the order $T = \mathcal{O}\left(\frac{1}{n\epsilon^2}\right)$. If gossip is used for approximate distributed averaging, and if the network is an expander, then the number of gossip iterations per round is also of the order $\mathcal{O}(\Theta(\log n))$, and so the speedup relative to using a single processor is nearly linear. This speedup is also observed in our experiments.

1.4 Paper Organization

The remainder of the paper is organized as follows. Section 2 introduces notation and states assumptions. Section 3 discusses exact and approximate protocols for distributed averaging. Section 5 describes the proposed algorithm and presents our main results. The proofs of all results are given in Section 6. Results of numerical experiments are presented in Section 7. We further discuss our results related to other work in the literature in Section 8, and we conclude in Section 9.

2 Notation and Assumptions

Throughout this paper, for a vector $x \in \mathbb{R}^d$, $\|x\|$ denotes the Euclidean (i.e., ℓ_2) norm, $\|x\| = x^T x$. The set of the first n natural numbers is denoted by $[n] \stackrel{\text{def}}{=} \{1, 2, \dots, n\}$. The d -dimensional vector with all entries equal to 0 (resp., all equal to 1) is written as $\mathbf{0}_d$ (resp. $\mathbf{1}_d$).

We focus on the two problem settings described in Sections 1.1 and 1.2 above. To re-iterate, the samples $x_i(t, s) \in \mathcal{X}$ arriving to the network are assumed to be independent and identically distributed, and they are drawn from an unknown distribution over \mathcal{X} , where $\mathcal{X} \subseteq \mathbb{R}^d$ is a d -dimensional Euclidean space for some $d \geq 1$. The objective is to minimize the regret $R_n(m)$ accumulated after m samples have been processed, in total, across the network. Our results and analysis below make use of the following assumptions.

Assumption 1. *The set \mathcal{W} is a closed, bounded, convex subset of \mathbb{R}^d with diameter $D = \max_{v, w \in \mathcal{W}} \|v - w\|$.*

Assumption 2. *The loss function $f(w, x)$ has the following properties.*

- a. (Convexity) *For all $x \in \mathcal{X}$, the loss function $f(w, x)$ is convex in w .*
- b. (Lipschitz continuity) *There exists a constant $L \geq 0$ such that for all $x \in \mathcal{X}$ and $w_1, w_2 \in \mathcal{W}$,*

$$|f(w_1, x) - f(w_2, x)| \leq L \|w_1 - w_2\|.$$

- c. (Lipschitz continuous gradients) *The loss function $f(w, x)$ is continuously differentiable with respect to w for all $x \in \mathcal{X}$, and there exists a constant $K \geq 0$ such that for all $x \in \mathcal{X}$ and $w_1, w_2 \in \mathcal{W}$,*

$$\|\nabla_w f(w_1, x) - \nabla_w f(w_2, x)\| \leq K \|w_1 - w_2\|.$$

Note that Assumptions 1 and 2 together imply that $f(w, x)$ also has bounded gradients; i.e., $\|\nabla_w f(w, x)\| \leq L$ for all $w \in \mathcal{W}$ and all $x \in \mathcal{X}$.

Let $F(w) = \mathbb{E}_x[f(w, x)]$. In each round, the algorithms we consider will make use of gradients $\nabla_w f(w, x)$. These can be viewed as stochastic or noisy versions of the gradient $\nabla_w F(w)$. The assumptions stated above (in particular, Lipschitz continuity of f and $\nabla_w f$) on the loss function are sufficient to ensure that the gradient operator commutes with the expectation [26], and thus $\nabla_w f(w, x)$ is an unbiased estimate for $\nabla_w F(w)$. In addition, we assume that the gradients have bounded variance. In particular:

Assumption 3. *There exists a constant $\sigma^2 \in (0, \infty)$ such that*

$$\mathbb{E}_x \left[\|\nabla_w f(w, x) - \nabla_w F(w)\|^2 \right] \leq \sigma^2.$$

After receiving data and suffering the loss in one round, node i may send/receive messages to/from other nodes before updating its predictor, $w_i(t+1)$, at the beginning of the next round. Each node is restricted to only communicate with its neighbours in a graph $G = (V, E)$ where $V = [n]$. Throughout this work we assume that:

Assumption 4. *The graph G is undirected and connected.*

Of course, G could be the complete graph (i.e., each node communicates with all the others). However, then the amount of communication required per node per round may not scale well with the size of the network, and so in general each node may only be neighbours with a few other nodes in the network.

3 Distributed Averaging Protocols

Before proceeding to the description of the algorithmic framework considered in this paper, we discuss distributed averaging protocols, which will be used as a mechanism to coordinate and synchronize the decision variables at different nodes.

Consider a connected network of n nodes where each node i holds a vector $y_i \in \mathbb{R}^d$. The goal of a distributed averaging protocol is for all nodes in the network to compute the average vector, $\bar{y} = \frac{1}{n} \sum_{i=1}^n y_i$. Let

$$y_i^+ = \text{DISTRIBUTEDAVERAGING}(y_i, i)$$

denote a function called at node i to invoke a distributed averaging protocol. The first argument is the vector, $y_i \in \mathbb{R}^d$, passed to the routine for averaging, and the second argument indicates this is the instance invoked at node i . The function returns a vector $y_i^+ \in \mathbb{R}^d$ as the output at node i .

In general, the output y_i^+ may not be identical at all nodes, and the quality of a distributed averaging protocol is measured through two metrics: the *accuracy* of the protocol and the *latency* of the protocol. A distributed averaging protocol is said to have accuracy $\delta \geq 0$ if it guarantees that, for all i ,

$$\|y_i^+ - \bar{y}\| \leq \delta$$

when the protocol returns. In the context of distributed online prediction, the latency of a distributed averaging protocol is measured with respect to the rate at which observations arrive to the network (and hence, the rate at which predictions must be made). A distributed averaging protocol is said to have latency $\mu > 0$ if a total of μ observations arrive to the network (i.e., μ/n observations at each node) during the time it takes for the protocol to return a value y_i^+ at each node. A (δ, μ) -distributed averaging protocol is one which guarantees that accuracy δ is achieved across the network while requiring latency at most μ to complete.

Typically the latency of a distributed averaging protocol is related to the size and structure of the network, as well as the accuracy required. In particular, we will be interested in how the latency scales both as a function of the accuracy required and as a function of the size of the network. Next we give examples of exact and approximate distributed averaging protocols.

3.1 Exact Distributed Averaging with ALLREDUCE

One way to exactly achieve distributed averaging in a network is to pass messages along a rooted spanning tree. Starting from the leaves of the spanning tree and assuming that all nodes know the size n of the network, nodes pass a message containing y_i/n to their parents in the tree. Once a node has received messages from all of its children, it computes the sum of the values received in these messages and also adds its own value, normalized by the size of the network, and then it passes a new message with this intermediate aggregate to its parent. This process repeats until the root node has received messages from all of its children. After the root sums the values received from its children together with its own (appropriately normalized) value, it holds the average \bar{y} . Then \bar{y} can be broadcast down the spanning tree to all nodes, and the protocol terminates when all nodes have received \bar{y} .

The operation just described corresponds to the ALLREDUCE primitive which is standard in existing message passing libraries for distributed computing [11]. It is clear that ALLREDUCE achieves an accuracy of $\delta = 0$, since all nodes hold the exact average when the protocol terminates. The latency of ALLREDUCE is proportional to the depth of the spanning tree and so, assuming G admits a balanced spanning tree where the branching factor that is bounded independent of n , the latency is of the order $\mu = \Theta(\log n)$.

Exact distributed averaging using ALLREDUCE is an attractive option when nodes communicate over a network which is stable and reliable. However, when nodes or links may fail, or when the communication channels between neighbouring nodes are unreliable, then there may be considerable overhead involved in establishing and maintaining a spanning tree across all nodes. In such a setting, it may be more desirable to use an approximate distributed averaging protocol with less overhead.

3.2 Approximate Distributed Averaging with Gossip Protocols

Gossip protocols are a family of simple message passing schemes that can be used for approximate distributed averaging. Here we focus our discussion on synchronous gossip protocols, but asynchronous versions exist and are equally usable within the algorithmic framework considered in this paper [3, 15, 36]; see [9, 21] for surveys and see [32] for a discussion of practical issues related to implementing asynchronous gossip protocols.

In a gossip protocol, rather than aggregating and broadcasting along a spanning tree, nodes iteratively pass messages to their neighbours in the communication graph G . In one iteration of a synchronous gossip algorithm, node i transmits a message to its neighbours containing its current value y_i , and it receives messages from each of its neighbours containing their values y_j . Then node i updates its value to be a convex combination of its previous values and those received from its neighbours. Specifically, let $P \in [0, 1]^{n \times n}$ denote a $n \times n$ matrix of update weights. We require that P respects the topology of the network in the

sense that $P_{i,j} > 0$ if and only if either $i = j$ or nodes i and j are neighbours in G (i.e., $(i, j) \in E$). After receiving messages from all of its neighbours, node i updates its value to

$$y_i \leftarrow \sum_{j=1}^n P_{i,j} y_j .$$

Note that since $P_{i,j} = 0$ if $(i, j) \notin E$ and $i \neq j$, the only non-zero terms in the summation are those corresponding to neighbours of i or to i itself, and so node i receives all the information it needs to compute the update via messages from its immediate neighbours. Let $y_i^{(0)}$ denote the initial vector at node i , and let

$$y_i^{(k)} = \text{GOSSIP}^k(y_i^{(0)}, i)$$

denote the vector output at node i after k gossip iterations. Then

$$y_i^{(k)} = \sum_{j=1}^n [P^k]_{i,j} y_j^{(0)} .$$

In addition to imposing that the sparsity structure of P matches the communication topology G , we assume that:

Assumption 5. *The matrix P is doubly stochastic; i.e., $\sum_{j=1}^n P_{i,j} = 1$ and $\sum_{i=1}^n P_{i,j} = 1$.*

This assumption, combined with the assumption that G is connected (Assumption 4) implies, via the Perron-Frobenius Theorem [27], that $\lim_{k \rightarrow \infty} P^k = \frac{1}{n} \mathbf{1}_n \mathbf{1}_n^T$. Clearly this is both necessary and sufficient to ensure that $y_i^{(k)} \rightarrow \bar{y}$ at every node i . Moreover, there are well-known bounds on the rate of convergence of these iterations. Since each iteration of gossip can be viewed as taking one discrete-time step of a diffusion over the network, it is not surprising that the rate of convergence depends on the communication topology through the second largest eigenvalue $\lambda_2(P)$ of the matrix P .

Lemma 1. *Let Assumptions 4 and 5 hold, let $\delta > 0$ be a given scalar, and let $y_i^{(k)} = \text{GOSSIP}^k(y_i^{(0)}, i)$ denote the output at node i after running k iterations of gossip for distributed averaging. Then it holds that*

$$\left\| y_i^{(k)} - \bar{y} \right\| \leq \delta$$

for all $i \in [n]$ if the number of gossip iterations k satisfies

$$k \geq \frac{\log \left(\frac{1}{\delta} \cdot 2\sqrt{n} \cdot \max_j \left\| y_j^{(0)} - \bar{y} \right\| \right)}{1 - \lambda_2(P)} . \quad (3)$$

The proof is given in Appendix A. The latency of gossip-based distributed averaging is proportional to the number of gossip iterations required, and it is evident from Lemma 1 that this depends on the size and structure of the network, both directly through \sqrt{n} as well as through $\lambda_2(P)$. It is well-known that the inverse spectral gap, $1/(1 - \lambda_2(P))$ is related to the structure of the network [6]. This relationship is especially important for understanding how the performance of gossip protocols (as well as the distributed optimization method introduced below) scale as a function of the network size and structure.

Of particular interest is the case when G is taken from the family of constant-degree expander graphs [25]. Then $1 - \lambda_2(P) = \Theta(1)$ as $n \rightarrow \infty$, and in addition, every node has the same number of neighbours. For gossip-based distributed averaging run over an expander graph or a complete graph, the resulting latency to obtain accuracy δ is

$$\mu = \mathcal{O} \left(\log \left(\frac{1}{\delta} \sqrt{n} \max_j \left\| y_j^{(0)} - \bar{y} \right\| \right) \right) .$$

In particular, if both δ and $\max_j \left\| y_j^{(0)} - \bar{y} \right\|$ are fixed constant as $n \rightarrow \infty$, then the latency scales as $\mathcal{O}(\log(n))$ which is exactly the same as for exact distributed averaging using ALLREDUCE.

In general, we will see that one can use any connected graph G satisfying the assumptions mentioned above. The expanders mentioned above are especially interesting because of the scaling properties mentioned. For less well-connected graphs, such as rings or two-dimensional grids, the inverse spectral gap $\frac{1}{1 - \lambda_2(P)}$ typically grows as a polynomial in n , and so more gossip iterations (and hence, more communications and a higher latency) are required to guarantee an accuracy of δ .

4 Distributed Dual Averaging with Approximate Mini-Batch Gradients

Next we present the general template for the distributed dual averaging algorithm using an approximate distributed averaging protocol to compute the mini-batch gradients. Node i maintains a primal variable $w_i(t) \in \mathbb{R}^d$ and a dual variable $z_i(t) \in \mathbb{R}^d$ which are both initialized to $w_i(1) = z_i(1) = \mathbf{0}_d$. In round $t \geq 1$, after receiving the samples $x_i(t, 1), \dots, x_i(t, \frac{b}{n})$ and suffering the corresponding loss $\sum_{s=1}^{\frac{b}{n}} f(w_i(t), x_i(t, s))$, node i computes the local average gradient

$$g_i(t) = \frac{n}{b} \sum_{s=1}^{\frac{b}{n}} \nabla_w f(w_i(t), x_i(t, s)) . \quad (4)$$

Then, the nodes update their dual variables by running an approximate distributed averaging protocol,

$$z_i(t+1) = \text{DISTRIBUTEDAVERAGING}(z_i(t) + g_i(t), i) . \quad (5)$$

Finally, each node updates its primal variable via the proximal projection,

$$w_i(t+1) = \arg \min_{w \in \mathcal{W}} \{ \langle w, z_i(t+1) \rangle + \beta(t+1)h(w) \} , \quad (6)$$

where $\{\beta(t)\}_{t=2}^\infty$ is a positive non-decreasing sequence of algorithm parameters and $h(w)$ is a 1-strongly convex proximal function; i.e., $h(0) = 0$, and for all $\theta \in [0, 1]$ and all $w, w' \in \mathcal{W}$,

$$h(\theta w + (1-\theta)w') \leq \theta \cdot h(w) + (1-\theta)h(w') - \frac{1}{2}\theta(1-\theta) \|w - w'\|^2 .$$

When working in Euclidean domains, as considered in this paper, the typical choice is to take $h(w) = \|w\|^2$; see [20, 29] for examples of other possible proximal functions that are more suitable in different domains.

For the framework just described, the same predictor $w_i(t)$ is used for all samples $x_i(t, s)$ processed at node i in round t . Consequently, the expression for the regret simplifies from (2) to

$$R_n(m) = \sum_{i=1}^n \sum_{t=1}^{\frac{m}{b+\mu}} \sum_{s=1}^{\frac{b+\mu}{n}} [f(w_i(t), x_i(t, s)) - f(w^*, x_i(t, s))] . \quad (7)$$

Note that, in the update (5), the argument passed to the distributed averaging protocol by node i is $z_i(t) + g_i(t)$, the sum of the dual variable and mini-batch gradient at node i . If an approximate distributed averaging protocol is used, then the output at node i is roughly

$$z_i(t+1) = \frac{1}{n} \sum_{i=1}^n (z_i(t) + g_i(t)) + \xi_i(t) \quad (8)$$

$$= \bar{z}(t) + \frac{1}{b} \sum_{i=1}^n \sum_{s=1}^{\frac{b}{n}} \nabla_w f(w_i(t), x_i(t, s)) + \xi_i(t) , \quad (9)$$

where $\xi_i(t)$ is an error term at node i and $\bar{z}(t)$ is the network-average dual variable at round t . If an approximate distributed averaging protocol, such as gossip, is used in (5) then an upper bound on the magnitude of the error $\xi_i(t)$ is guaranteed and can be made arbitrarily small, typically at the cost of increased communications. Thus, each node effectively updates the dual variable by adding the mini-batch gradient using the b gradient samples obtained across the network.

Before proceeding to our main results, we remark that other distributed versions of the dual averaging algorithm using approximate distributed averaging have proposed in the literature [10]. There are two differences between the algorithm considered in [10] and the framework described above. The first is that the algorithm of [10] only allows for a single gossip iteration in each round, whereas the framework above

allows for (5) to involve multiple gossip iterations. The second, and more subtle point is that the algorithm of [10] only averages the dual variables $z_i(t)$, which would be equivalent to having node i call

$$z_i(t+1) = \text{DISTRIBUTEDAVERAGING}(z_i(t), i) + g_i(t) ,$$

from which it is evident that the $z_i(t)$ no longer benefits from the same mini-batch gradient updates as in 9; rather, averaging of the mini-batch of gradients $g_i(t)$ from round t is delayed by a round and only appears in $z_i(t+2)$.

5 Main Results

5.1 Online Prediction

5.1.1 Regret Bound for General Approximate Distributed Averaging Protocol

Our first main result demonstrates that the optimal $\mathcal{O}(\sqrt{m})$ regret bound can be achieved despite using approximate distributed averaging protocols in the distributed mini-batch framework.

Theorem 1. *Let Assumptions 1–4 hold. Let $z_i(t)$ and $w_i(t)$ be the values generated at node i by the algorithm (4)–(6) described above for $t \geq 1$, and with $\beta(t) = K + \sqrt{\frac{t}{b+\mu}}$. Suppose that the distributed averaging protocol used in (5) has accuracy δ and latency μ . Then the expected regret after processing m samples in total across the network satisfies the bound*

$$\begin{aligned} \mathbb{E}[R_n(m)] &\leq (b + \mu)[F(\bar{w}(1)) - F(w^*) + Kh(w^*)] + \frac{3}{4}\delta^2 K^2 (b + \mu)^{5/2} \\ &\quad + \left[2\sigma^2 \frac{b + \mu}{b} + 2\delta K D \frac{b + \mu}{n} + 2\delta L(b + \mu) \right] \sqrt{m} . \end{aligned}$$

The proof is given in Section 6.1. The result above bounds the expected regret as a function of the number of samples m processed by the network, the mini-batch size b , the accuracy δ and latency μ of the distributed averaging protocol, and other problem-dependent constants. If the accuracy of the distributed averaging protocol is sufficiently small relative to its latency, then we obtain the following simplified bound.

Corollary 1. *If $\delta \leq \frac{1}{b+\mu}$ then $\mathbb{E}[R_n(m)] = \mathcal{O}(b + \sqrt{m})$. Moreover, if $b = m^\rho$ for a constant $\rho \in (0, 1/2)$, then $\mathbb{E}[R_n(m)] = \mathcal{O}(\sqrt{m})$.*

From the corollary, we see that to obtain the optimal regret scaling we must use a distributed averaging protocol for which the accuracy and latency satisfy the relationship $\delta \leq \frac{1}{b+\mu}$. Then, the optimal rate of $\mathcal{O}(\sqrt{m})$ follows by scaling b appropriately with respect to the total number of data samples m are to be processed.

5.1.2 Approximate Mini-Batch Gradients via Gossip-based Distributed Averaging

Next we investigate conditions under which gossip-based distributed averaging achieves the optimal rates in terms of expected regret. Recall that the gossip protocol for distributed averaging described in Section 3.2 is guaranteed to provide δ -accurate output at each node, in the sense that $\|y_j^{(k)} - \bar{y}\| \leq \delta$ for all nodes $j \in [n]$, if the number of iterations k satisfies

$$k \geq \frac{\log\left(\frac{1}{\delta} \cdot 2\sqrt{n} \cdot \max_j \|y_j^{(0)} - \bar{y}\|\right)}{1 - \lambda_2(P)} ,$$

where $y_j^{(0)} \in \mathbb{R}^d$ is the initial vector at node j , and \bar{y} is the average of the initial vectors.

The latency of gossip is proportional to the number of iterations executed; hence, let us suppose that there exists a positive integer γ such that $\mu = \gamma k$; i.e., γ is the number of samples arriving to the system during a single gossip iteration.

From Corollary 1 above, we know that the optimal regret bound is obtained if the accuracy and delay satisfy the relationship $\delta \leq \frac{1}{b+\mu} = \frac{1}{b+\gamma k}$. Intuitively we expect that it should be feasible to achieve such a bound using gossip iterations on any graph since the accuracy δ decays exponentially in k (see Lemma 1) whereas the delay μ is linear in k . This intuition is made precise in the following theorem.

Theorem 2. *Let Assumptions 1–5 hold. Let $z_i(t)$ and $w_i(t)$ denote the sequence of values generated at node i by the algorithm (4)–(6) described above for $t \geq 1$. Suppose that the distributed averaging operation in step (5) is implemented by running the gossip protocol described in Section 3.2 for $k > 0$ iterations; i.e., for step (5) take*

$$z_i(t+1) = \text{GOSSIP}^k(z_i(t) + g_i(t), i) .$$

Suppose, in addition, that there exists a constant $\gamma > 0$ independent of k such that the latency of the gossip protocol is $\mu = \gamma k$, and assume that $\gamma < b$. Take the algorithm parameters to be $\beta(t) = K + \sqrt{\frac{t}{b+\gamma k}}$. If the number of gossip iterations k executed at each round of the algorithm satisfies

$$k \geq \left\lceil \left(\frac{\log(4Lb\sqrt{n}) + \log(\frac{1}{1-\lambda_2(P)})}{1-\lambda_2(P)} + \frac{1}{2Lb} + 1 \right) \frac{1}{1-\frac{\gamma}{b}} \right\rceil , \quad (10)$$

then, for all $j \in [n]$ and all $t \geq 0$, the values $z_j(t+1)$ satisfy

$$\|z_j(t+1) - \bar{z}(t+1)\| \leq \frac{1}{b+\mu} ,$$

where $\bar{z}(t+1) = \frac{1}{n} \sum_{j=1}^n z_j(t+1)$. Thus, the expected regret is bounded as

$$\mathbb{E}[R_n(m)] = \mathcal{O}(b + \sqrt{m}) .$$

Moreover, if the mini-batch size is chosen so that $b = m^\rho$ for some $\rho \in (0, 1/2)$ then $\mathbb{E}[R_n(m)] = \mathcal{O}(\sqrt{m})$.

The proof is given in Section 6.3. The assumption that $\gamma < b$ simply states that the number of samples γ arriving to the network per gossip iteration is less than the total size of the mini-batch. If this is not the case, then the number of samples being lost during the communication phase of the network is larger than the number of samples used in the algorithm updates (4)–(6), which is not very sensible. For cases of practical interest, one expects that the total latency, μ , and hence γ , will be much smaller than b (e.g., γ remains constant as m grows, whereas b grows as a function of m) [8].

Theorem 2 provides a condition on the number of gossip iterations k executed in each round in order to guarantee that the optimal expected regret bound is achieved. From the perspective of scaling laws (as the network size n grows), the three parameters which may depend on n are the inverse spectral gap $1/(1-\lambda_2(P))$, the proportionality constant γ , and the mini-batch size b . The way in which the inverse spectral gap depends on n is related to structural properties of the network topology, as discussed in Section 3.2 above.

The proportionality constant γ is the number of samples being processed across the network during each gossip iteration (i.e., the latency of each iteration), and it is related to the amount of time it takes to communicate for each gossip iteration. This depends on both the type of communication mechanism available and possibly the structure of the network topology. For example, if nodes communicate using point-to-point messages, then in each iteration, every node must send and receive a message from each of its neighbours, and thus γ is proportional to the maximum degree of the network. Hence, $\gamma = \mathcal{O}(n)$; in the most extreme case, each node sends a message to every other node in each iteration of gossip, and so $\gamma \propto n$. In more practical cases we expect that $\gamma = o(n)$; for example, if the maximum degree in the network grows logarithmically, as it does with high probability in many common random graph models (including the Erdős-Renyi model and the random geometric graph model), then $\gamma = \mathcal{O}(\log n)$. If the degree of each node remains fixed as n grows, as in the constant-degree expander model [25]), then $\gamma = \mathcal{O}(1)$. It also holds that $\gamma = o(n)$ if nodes communicate using local broadcasts (e.g., over a reliable wireless channel).

The mini-batch size b is linearly proportional to n since b is the total number of samples processed across the network in each round (and so the number of samples per node is b/n). Since γ only enters the expression (10) via the ratio γ/b , and since $\gamma = \mathcal{O}(n)$, we see that the term $1/(1-\frac{\gamma}{b}) = \mathcal{O}(1)$; when $\gamma = \Theta(n)$, then $1/(1-\frac{\gamma}{b})$ remains a constant less than one, and when $\gamma = o(n)$ then $1/(1-\frac{\gamma}{b})$ converges to

1 as $n \rightarrow \infty$. Thus, although the particular communication mechanism affects the constants (through γ), it does not affect the overall rate at which k grows as a function of n .

Fixing a particular class of network topology determines the relevant laws for $1/(1 - \lambda_2(P))$, γ , and b , and the communication scaling law (how the amount of communication per round, k , grows as a function of n) can be determined. Again, of particular interest is the case where G is a constant-degree expander graph. Then $\gamma = o(n)$ because the maximum degree of the graph is a constant independent of n , and $1/(1 - \lambda_2(P)) = \Theta(1)$. Thus, for G being a constant-degree expander graph, $k = \Theta(\log n)$ iterations of gossip suffices.

5.2 Stochastic Optimization

5.2.1 General Optimization Error Bound

Next, we turn to the problem of stochastic optimization. Recall that in distributed stochastic optimization, a network of n nodes aims to cooperatively solve the problem

$$\begin{aligned} & \text{minimize} && F(w) \stackrel{\text{def}}{=} \mathbb{E}_x[f(w, x)] \\ & \text{subject to} && w \in \mathcal{W} . \end{aligned}$$

We assume a similar framework to that considered above, where $f(w, x)$ and \mathcal{W} satisfy Assumptions 1–3, and where nodes communicate over a network satisfying Assumption 4.

Unlike distributed online prediction, new samples do not arrive while the nodes are communicating, and so no loss is suffered (in terms of the optimization error) due to latency. On the other hand, since the motivation of using a distributed system to solve the problem is to *quickly* obtain an accurate solution, it is of interest to know how the time to reach an ϵ -accurate solution scales with the number of nodes in the network.

We consider a method which operate in a similar manner to those described in Section 4 for distributed online prediction. Node i maintains its own copy of the optimization variable $w_i(t)$, and the values at each node are updated in synchronous rounds. Within each round, nodes may process a mini-batch of samples locally, communicate with their neighbours, and then update their local variables.

For stochastic optimization, we augment the dual averaging algorithm described in Section 4. After performing the steps (4)–(6), node i also computes the running average predictor,

$$\hat{w}_i(t+1) \stackrel{\text{def}}{=} \frac{1}{t+1} \sum_{t'=1}^{t+1} w_i(t') , \tag{11}$$

In stochastic optimization, performance is measured in terms of the optimality gap. After the network has completed T rounds, the optimality gap is given by

$$\Delta_n(T) \stackrel{\text{def}}{=} \max_{i \in [n]} \mathbb{E}[F(\hat{w}_i(T))] - F(w^*) .$$

It is well-known that there is a direct connection between the performance of online prediction methods and stochastic optimization in the single-processor setting [35], where Jensen’s inequality gives that

$$\Delta_1(T) \leq \frac{1}{T} \mathbb{E}[R_1(T)] .$$

Thus, any online prediction algorithm with expected regret scaling as $\mathbb{E}[R_1(m)] = \mathcal{O}(\sqrt{m})$ can be used to obtain an optimization error of $\Delta_1(T) = \mathcal{O}\left(\frac{1}{\sqrt{T}}\right)$. Equivalently, in order to guarantee that $\Delta_1(T) \leq \epsilon$ using such an algorithm, we need to run at least $T \propto \frac{1}{\epsilon^2}$ rounds. In the distributed setting, we obtain a similar bound for the distributed dual averaging algorithm with approximate mini-batches.³

³Note that, unlike in the single-processor setting just discussed, we cannot simply apply Jensen’s inequality to the regret bound from Theorem 1 to obtain this result. This is because the different processors produce sequences $\{\hat{w}_i(t)\}$ which are not exactly the same due to the use of an approximate distributed averaging protocol, and these errors must be properly accounted for to obtain a bound on the optimization error.

Theorem 3. *Let Assumptions 1–4 hold. Let $\{\widehat{w}_i(t)\}_{t \geq 1}$ denote the values generated at node i by repeating the iterations (4)–(6) and (11) with algorithm parameters $\beta(t) = K + \sqrt{\frac{t}{b}}$. For some positive integer $C \geq 1$, set the mini-batch size to $b = Cn$. Assume that the approximate distributed averaging protocol used in (5) guarantees accuracy $\delta \leq \frac{1}{b}$ in each round. Then after T rounds, we have*

$$\Delta_n(T) \leq \frac{1}{T} \left[F(0) - F(w^*) + Kh(w^*) + \frac{3K^2}{4\sqrt{Cn}} \right] + \frac{1}{\sqrt{T}} \left[\frac{\sigma^2}{4\sqrt{Cn}} + \frac{2KD}{n\sqrt{Cn}} + \frac{2L}{\sqrt{Cn}} \right]. \quad (12)$$

The proof is given in Section 6.4. The result above shows that, for a fixed network size n , the optimality gap after T rounds is of the order $\Delta_n(T) = \mathcal{O}\left(\frac{1}{\sqrt{nT}}\right)$. Thus, for $\epsilon > 0$, in order to guarantee that $\Delta_n(T) \leq \epsilon$ we need to perform at least $T \propto \frac{1}{n\epsilon^2}$ rounds. Similar to the serial setting, the number of rounds scales as $\frac{1}{\epsilon^2}$. However, importantly, we see that in the distributed setting, the number of rounds is reduced by a factor of n , and so we obtain a linear scale-up in terms of the network size. In order to translate this scaling in the number of rounds T to a speedup in terms of the runtime, we also need to account for the time spent communicating in each round relative to the time spent computing. This is the topic of the next subsection.

5.2.2 Latency and Scaling

In distributed stochastic optimization we do not consider latency in the definition of the optimization error $\Delta_n(T)$ because, unlike the online prediction setting where inputs arrive in a streaming manner and predictions must be made in real-time,⁴ in stochastic optimization no additional loss or error needs to be accrued while the network is communicating. The primary motivation for using a distributed system in the context of stochastic optimization is to achieve accurate results *faster* than with a single node, by exploiting the ability to process inputs in parallel.

Theorem 3 above shows that we achieve a linear speedup in the number of iterations required to guarantee an optimization error of no more than ϵ . Specifically, $T \propto \frac{1}{n\epsilon^2}$. However, we should also account for the latency involved in communication in order to model the overall runtime to reach accuracy ϵ . To this end, we adapt the approach of [31] for modelling runtimes of distributed optimization methods.

Let us assume that time is normalized so that it takes a processor 1 unit of time to process a single sample x (i.e., to evaluate $\nabla_w f(w, x)$ once), and let $\tau > 0$ denote the number of time units it takes a node to transmit its dual variable to one neighbour.⁵ In the serial (single-processor) setting, the dual averaging algorithm must process $T \propto \frac{1}{\epsilon^2}$ samples in order to guarantee an optimality gap of the order $\Delta_1(T) = \mathcal{O}(\epsilon)$, and so clearly the runtime also scales as $\mathcal{O}\left(\frac{1}{\epsilon^2}\right)$ time units.

Suppose we run a gossip protocol for k iterations to implement the approximate distributed averaging operation (5) in every round. In each iteration, each node must send a copy of its dual variable to each of its neighbours, so the time spent communicating per round is $\tau \cdot k \cdot \deg(G)$, where $\deg(G)$ denotes the maximum node degree in G . Before communicating, the nodes each process b/n samples, in parallel, to compute $g_i(t)$, so the total time per round is $\frac{b}{n} + \tau \cdot k \cdot \deg(G)$ time units. Combining this with the fact that the distributed algorithm needs to run for $T \propto \frac{1}{n\epsilon^2}$ rounds, we have that the overall runtime is proportional to

$$\frac{1}{n\epsilon^2} \left(\frac{b}{n} + \tau \cdot k \cdot \deg(G) \right) = \frac{1}{n\epsilon^2} (C + \tau \cdot k \cdot \deg(G)) \quad \text{time units,}$$

where k is the number of gossip iterations required to guarantee $\delta \leq \frac{1}{b}$ in every round.

To determine the appropriate value of k we can directly apply Lemma 1 in this case to find that $\delta \leq \frac{1}{b}$ if we run at least

$$k \geq \frac{\log \left(b \cdot 2\sqrt{n} \cdot \max_j \|y_j^{(0)} - \bar{y}\| \right)}{1 - \lambda_2(P)}$$

⁴Recall that the motivation for using a distributed system for online prediction was to enable handling higher rates of predictions.

⁵We will disregard the time it takes to compute the proximal projection (6); relative to the time spent aggregating gradients in (4) and the time spent communicating in (5), which both depend on the mini-batch size b and network size n , the time to compute (6) is $\mathcal{O}(1)$. Often, there is a closed form expression for the solution to (6).

gossip iterations, where $y_j^{(0)} = z_j(t) + g_j(t)$ is the initial value at each node when the gossip protocol is invoked to compute $z_j(t+1)$. Note that

$$\begin{aligned} \|z_j(t) + g_j(t) - (\bar{z}(t) + \bar{g}(t))\| &\leq \|z_j(t) - \bar{z}(t)\| + \|g_j(t) - \bar{g}(t)\| \\ &\leq \|z_j(t) - \bar{z}(t)\| + 2L, \end{aligned}$$

where the second inequality follows after a second application of the triangle inequality along with recalling that Assumptions 1 and 2 imply that the gradients $g_j(t)$ are bounded (and hence $\bar{g}(t)$ is too). In addition, we claim that $\|z_j(t) - \bar{z}(t)\| \leq \frac{1}{b}$ for all t . Clearly this holds for the first round since $z_j(1) = \mathbf{0}_d$ at every node. Proceeding inductively, suppose that $\|z_j(t) - \bar{z}(t)\| \leq \frac{1}{b}$. Then choosing

$$k \geq \frac{\log(b \cdot 2\sqrt{n} \cdot (\frac{1}{b} + 2L))}{1 - \lambda_2(P)} \propto \frac{\log(n)}{1 - \lambda_2(P)}$$

ensures that $\|z_j(t+1) - \bar{z}(t+1)\| \leq \frac{1}{b}$, and so it holds for all $t \geq 1$.

Returning to the runtime scaling, we have that overall runtime when gossip is used for approximate distributed averaging is proportional to

$$\frac{1}{n\epsilon^2} \left(C + \frac{\tau \log(n)}{1 - \lambda_2(P)} \cdot \deg(G) \right) \quad \text{time units.}$$

When G is taken from the family of constant-degree expander graphs then $\deg(G) = \Theta(1)$ and $\frac{1}{1 - \lambda_2(P)} = \Theta(1)$, and so the overall runtime scales as

$$\frac{1}{n\epsilon^2} (C + \tau \log(n)) \quad \text{time units.}$$

Thus, the runtime scales nearly linearly, as $\mathcal{O}\left(\frac{\log n}{n}\right)$, when communication is over an expander graph. Note that this is identical to the scaling achieved when exact distributed averaging is used, since aggregating along a spanning tree also takes $\mathcal{O}(\log n)$ time units.

6 Analysis

6.1 Proof of Theorem 1

Let

$$X(t) \stackrel{\text{def}}{=} \left\{ x_i(t', s) : 1 \leq i \leq n, 1 \leq t' \leq t \text{ and } 1 \leq s \leq \frac{b+\mu}{n} \right\}$$

denote the set of all data processed across the network during the first t rounds. Observe that the predictor $w_i(t)$ computed using (6) is independent of $x_i(t, s)$ given $X(t-1)$. Recall that $\mathbb{E}[\cdot]$ denotes the expectation with respect to the entire data sequence $X(\frac{m}{b+\mu})$, i.e., with respect to the entire collection of m samples processed by the network when computing $\mathbb{E}[R_n(m)]$. Conditioning on $X(t-1)$ gives

$$\begin{aligned} \mathbb{E}[f(w_i(t), x_i(t, s))] &= \mathbb{E}[\mathbb{E}[f(w_i(t), x_i(t, s)) | X(t-1)]] \\ &= \mathbb{E}[F(w_i(t))]. \end{aligned}$$

Also, clearly $\mathbb{E}[f(w^*, x_i(t, s))] = \mathbb{E}[F(w^*)]$. Therefore, for $t \leq \frac{m}{b+\mu}$,

$$\mathbb{E}[R_n(m)] = \sum_{t=1}^{\frac{m}{b+\mu}} \sum_{i=1}^n \sum_{s=1}^{\frac{b+\mu}{n}} [\mathbb{E}[F(w_i(t)) - F(w^*)]], \quad (13)$$

and we can focus on bounding the performance with respect to the expected loss function, $F(w)$.

Each predictor $w_i(t)$ is a proximal projection of the dual variable $z_i(t)$ as given in (6). In the sequel we will use the fact that this projection is Lipschitz continuous, as given by the following lemma. (See, e.g., [10] for a proof.)

Lemma 2. Let $h: \mathbb{R}^d \rightarrow \mathbb{R}$ be a strongly convex function and consider the proximal projection

$$\Pi_{\mathcal{W}}(z, \beta) \stackrel{\text{def}}{=} \arg \min_{w \in \mathcal{W}} \{ \langle w, z \rangle + \beta h(w) \} .$$

For any vectors $z, z' \in \mathbb{R}^d$ and for any positive scalar $\beta > 0$,

$$\|\Pi_{\mathcal{W}}(z, \beta) - \Pi_{\mathcal{W}}(z', \beta)\| \leq \frac{1}{\beta} \|z - z'\| .$$

By assumption, the distributed averaging algorithm used in (5) produces an output $z_i(t+1)$ at node i which satisfies $\|z_i(t+1) - \bar{z}(t+1)\| \leq \delta$ for all i , where

$$\bar{z}(t+1) \stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^n (z_i(t) + g_i(t)) . \quad (14)$$

If all nodes were to exactly compute $\bar{z}(t+1)$ and use this in the projection step (6) then they would obtain the predictor

$$\bar{w}(t+1) \stackrel{\text{def}}{=} \arg \min_{w \in \mathcal{W}} \{ \langle w, \bar{z}(t+1) \rangle + \beta(t+1)h(w) \} . \quad (15)$$

To bound the expected regret across the network in the general case where an approximate distributed averaging protocol is used, we relate the expected regret to the regret that would be accumulated using the reference predictor sequence, $\{\bar{w}(t)\}_{t \geq 1}$, and then we bound the deviation between the performance of the predictors $\{w_i(t)\}_{t \geq 1}$ from that of the reference predictors $\{\bar{w}(t)\}_{t \geq 1}$.

Adding and subtracting $F(\bar{w}(t))$ in (13) gives

$$\begin{aligned} \mathbb{E}[R_n(m)] &= \mathbb{E} \left[\sum_{t=1}^{\frac{m}{b+\mu}} \sum_{i=1}^n \sum_{s=1}^{\frac{b+\mu}{n}} [F(\bar{w}(t)) - F(w^*) + F(w_i(t)) - F(\bar{w}(t))] \right] \\ &= \mathbb{E} \left[(b+\mu) \sum_{t=1}^{\frac{m}{b+\mu}} [F(\bar{w}(t)) - F(w^*)] + \frac{b+\mu}{n} \sum_{t=1}^{\frac{m}{b+\mu}} \sum_{i=1}^n [F(w_i(t)) - F(\bar{w}(t))] \right] . \end{aligned}$$

Since $f(w, x)$ is Lipschitz continuous in w for all x by Assumption 2, it follows that $F(w)$ is Lipschitz continuous with the same constant. Using this property along with Lemma 2 and linearity of expectation gives

$$\begin{aligned} \mathbb{E}[R_n(m)] &\leq \mathbb{E} \left[(b+\mu) \sum_{t=1}^{\frac{m}{b+\mu}} [F(\bar{w}(t)) - F(w^*)] + \frac{(b+\mu)L}{n} \sum_{t=1}^{\frac{m}{b+\mu}} \sum_{i=1}^n \|w_i(t) - \bar{w}(t)\| \right] \\ &\leq (b+\mu) \sum_{t=1}^{\frac{m}{b+\mu}} \mathbb{E}[F(\bar{w}(t)) - F(w^*)] + \frac{(b+\mu)L}{n} \sum_{t=1}^{\frac{m}{b+\mu}} \sum_{i=1}^n \frac{1}{\beta(t)} \mathbb{E}[\|z_i(t) - \bar{z}(t)\|] \\ &\leq (b+\mu) \sum_{t=1}^{\frac{m}{b+\mu}} \mathbb{E}[F(\bar{w}(t)) - F(w^*)] + (b+\mu)L\delta \sum_{t=1}^{\frac{m}{b+\mu}} \frac{1}{\beta(t)} , \end{aligned} \quad (16)$$

where the last inequality follows based on the assumed accuracy of the approximate distributed averaging protocol. The main work remaining to prove Theorem 1 involves bounding the first sum in (16). To this end we have:

Lemma 3. Let the assumptions of Theorem 1 hold, and let $a(t) = \sqrt{\frac{t}{b+\mu}}$. Then

$$\begin{aligned} \sum_{t=1}^{\frac{m}{b+\mu}} \mathbb{E}[F(\bar{w}(t)) - F(w^*)] &\leq F(\bar{w}(1)) - F(w^*) + (K + a(\frac{m}{b+\mu}))h(w^*) + \frac{\sigma^2}{4b} \sum_{t=1}^{\frac{m}{b+\mu}-1} \frac{1}{a(t)} \\ &\quad + \frac{\delta^2 K^2}{4} \sum_{t=1}^{\frac{m}{b+\mu}-1} \frac{1}{a(t)(K + a(t))^2} + \frac{\delta K D}{n} \sum_{t=1}^{\frac{m}{b+\mu}-1} \frac{1}{K + a(t)} . \end{aligned}$$

The proof of Lemma 3 is given in Section 6.2. This addresses the first sum in (16). In order to bound the remaining terms, with $a(t) = \sqrt{\frac{t}{b+\mu}}$ as in the statement of the lemma, note that

$$\begin{aligned} \sum_{t=1}^{\frac{m}{b+\mu}-1} \frac{1}{\beta(t)} &= \sum_{t=1}^{\frac{m}{b+\mu}-1} \frac{1}{K + a(t)} \\ &\leq \sum_{t=1}^{\frac{m}{b+\mu}-1} \frac{1}{a(t)} \\ &\leq \sqrt{b+\mu} \left(1 + \int_1^{\frac{m}{b+\mu}} t^{-1/2} dt \right) \\ &\leq 2\sqrt{m}, \end{aligned}$$

and

$$\begin{aligned} \sum_{t=1}^{\frac{m}{b+\mu}-1} \frac{1}{a(t)(K + a(t))^2} &\leq \sum_{t=1}^{\frac{m}{b+\mu}-1} \frac{1}{a(t)^3} \\ &= (b+\mu)^{3/2} \sum_{t=1}^{\frac{m}{b+\mu}} t^{-3/2} \\ &\leq (b+\mu)^{3/2} \left(1 + \int_1^{\frac{m}{b+\mu}} t^{-3/2} dt \right) \\ &\leq 3(b+\mu)^{3/2}. \end{aligned}$$

Combining these inequalities with Lemma 3 and (16) gives

$$\begin{aligned} \mathbb{E}[R_n(m)] &\leq (b+\mu) \left[F(\bar{w}(1)) - F(w^*) + \left(K + \frac{\sqrt{m}}{b+\mu} \right) h(w^*) + \frac{2\sigma^2}{b} \sqrt{m} + \frac{3}{4} \delta^2 K^2 (b+\mu)^{3/2} + \frac{2\delta K D}{n} \sqrt{m} \right] \\ &\quad + 2\delta(b+\mu)L\sqrt{m} \\ &= (b+\mu) [F(\bar{w}(1)) - F(w^*) + Kh(w^*)] + \frac{3}{4} \delta^2 K^2 (b+\mu)^{5/2} \\ &\quad + \left[2\sigma^2 \frac{b+\mu}{b} + 2\delta K D \frac{b+\mu}{n} + 2\delta L(b+\mu) \right] \sqrt{m}, \end{aligned}$$

which completes the proof of Theorem 1.

6.2 Proof of Lemma 3

We seek to upper-bound the sum

$$\sum_{t=1}^{\frac{m}{b+\mu}} \mathbb{E}[F(\bar{w}(t)) - F(w^*)],$$

the regret of the averaged sequence of predictors $\{\bar{w}(t)\}$ after $\frac{m}{b+\mu}$ rounds. From the definition of $\bar{z}(t+1)$ in (14) and from (5), we have

$$\bar{z}(t+1) = \bar{z}(t) + \bar{g}(t),$$

where

$$\begin{aligned} \bar{g}(t) &= \frac{1}{n} \sum_{i=1}^n g_i(t) \\ &= \frac{1}{b} \sum_{i=1}^n \sum_{s=1}^{\frac{b}{n}} \nabla_w f(w_i(t), x_i(t, s)). \end{aligned}$$

Note that, for $\bar{w}(t)$ as defined in (15), $\bar{w}(t) \neq \frac{1}{n} \sum_{i=1}^n w_i(t)$ since the projection operator is not linear. Note also that the sequence $\bar{w}(t)$ is not equivalent to running a serial (i.e., single processor) version of dual averaging with mini-batches since the gradients $g_i(t)$ computed at each node (cf., (4)) are evaluated at different points, $w_i(t)$, whereas a serial implementation would use

$$\hat{g}(t) = \frac{1}{b} \sum_{i=1}^n \sum_{s=1}^{\frac{b}{n}} \nabla_w f(\bar{w}(t), x_i(t, s)) ,$$

in place of $\bar{g}(t)$, with all gradients evaluated at $\bar{w}(t)$.

There are two sources of errors in the gradients which must be taken into account. The first source stems from the observation just mentioned, that the updates at different nodes involve evaluating gradients at different locations. The second source of errors is due to the randomness of the gradients $\nabla_w f(w, x)$ calculated using data x rather than directly using the gradients $\nabla_w F(w)$. Let us define two gradient error vectors,

$$\begin{aligned} q(t) &\stackrel{\text{def}}{=} \hat{g}(t) - \nabla_w F(\bar{w}(t)), \text{ and} \\ r(t) &\stackrel{\text{def}}{=} \bar{g}(t) - \hat{g}(t) . \end{aligned}$$

These error vectors satisfy the properties stated in the following lemma that is proved in Appendix B.

Lemma 4. *Let Assumptions 1–3 hold. Then for $t = 1, \dots, \frac{m}{b+\mu}$,*

$$\mathbb{E}[\langle q(t), w^* - \bar{w}(t) \rangle] = 0, \quad (17)$$

$$\mathbb{E}[\langle r(t), w^* - \bar{w}(t) \rangle] \leq \frac{KD}{n\beta(t)} \sum_{i=1}^n \mathbb{E}[\|z_i(t) - \bar{z}(t)\|], \quad (18)$$

$$\mathbb{E}[\|q(t)\|^2] \leq \frac{\sigma^2}{b}, \quad (19)$$

$$\mathbb{E}[\|r(t)\|^2] \leq \frac{K^2}{n^2\beta(t)^2} \sum_{i=1}^n \sum_{j=1}^n \mathbb{E}[\|z_i(t) - \bar{z}(t)\| \|z_j(t) - \bar{z}(t)\|], \quad (20)$$

where $\mathbb{E}[\cdot]$ denotes the expectation with respect to the joint distribution over the set of random variables $X(\frac{m}{b+\mu})$.

Let us denote the linear (first-order) model for $F(w)$ evaluated at $\bar{w}(t)$ by

$$l_t(w) \stackrel{\text{def}}{=} F(\bar{w}(t)) + \langle \nabla_w F(\bar{w}(t)), w - \bar{w}(t) \rangle .$$

We also make use of an approximation to $l_t(w)$ where $\nabla_w F(\bar{w}(t))$ is replaced with $\bar{g}(t)$,

$$\begin{aligned} \hat{l}_t(w) &\stackrel{\text{def}}{=} F(\bar{w}(t)) + \langle \bar{g}(t), w - \bar{w}(t) \rangle \\ &= l_t(w) + \langle q(t), w - \bar{w}(t) \rangle + \langle r(t), w - \bar{w}(t) \rangle . \end{aligned} \quad (21)$$

From Assumption 2, we obtain that $F(w)$ is convex and has Lipschitz continuous gradients. It follows that (see, e.g., Lemma 1.2.3 in [19])

$$\begin{aligned} F(\bar{w}(t+1)) &\leq l_t(\bar{w}(t+1)) + \frac{K}{2} \|\bar{w}(t+1) - \bar{w}(t)\|^2 \\ &= \hat{l}_t(\bar{w}(t+1)) - \langle q(t), \bar{w}(t+1) - \bar{w}(t) \rangle - \langle r(t), \bar{w}(t+1) - \bar{w}(t) \rangle + \frac{K}{2} \|\bar{w}(t+1) - \bar{w}(t)\|^2 . \end{aligned}$$

Let $a(t) = \beta(t) - K$. Applying the Cauchy-Schwarz inequality and adding and subtracting $a(t)$ in the term

multiplying $\|\bar{w}(t+1) - \bar{w}(t)\|^2$ gives

$$\begin{aligned}
F(\bar{w}(t+1)) &\leq l_t(\bar{w}(t+1)) + \frac{K + a(t)}{2} \|\bar{w}(t+1) - \bar{w}(t)\|^2 \\
&\quad + \|q(t)\| \|\bar{w}(t+1) - \bar{w}(t)\| - \frac{a(t)}{4} \|\bar{w}(t+1) - \bar{w}(t)\|^2 \\
&\quad + \|r(t)\| \|\bar{w}(t+1) - \bar{w}(t)\| - \frac{a(t)}{4} \|\bar{w}(t+1) - \bar{w}(t)\|^2 .
\end{aligned} \tag{22}$$

Observe that

$$\begin{aligned}
&\|q(t)\| \|\bar{w}(t+1) - \bar{w}(t)\| - \frac{a(t)}{4} \|\bar{w}(t+1) - \bar{w}(t)\|^2 \\
&= \frac{\|q(t)\|^2}{4a(t)} - \left[\frac{\|q(t)\|}{\sqrt{4a(t)}} - \sqrt{\frac{a(t)}{4}} \|\bar{w}(t+1) - \bar{w}(t)\| \right]^2 \\
&\leq \frac{\|q(t)\|^2}{4a(t)} ,
\end{aligned}$$

and a similar bound holds if $q(t)$ is replaced by $r(t)$.

Using these bounds in (22) gives

$$\begin{aligned}
F(\bar{w}(t+1)) &\leq \hat{l}_t(\bar{w}(t+1)) + \frac{K + a(t)}{2} \|\bar{w}(t+1) - \bar{w}(t)\|^2 \\
&\quad + \frac{\|q(t)\|^2}{4a(t)} + \frac{\|r(t)\|^2}{4a(t)} .
\end{aligned} \tag{23}$$

Next, we make use of two lemmas to simplify this expression further. The first lemma gives a relation between $\bar{w}(t)$ and $\hat{l}_t(\bar{w}(t))$.

Lemma 5. *Let $\bar{z}(t)$ and $\bar{w}(t)$ be defined as in (14) and (15). Then*

$$\bar{w}(t) = \arg \min_{w \in \mathcal{W}} \left\{ \sum_{t'=1}^{t-1} \hat{l}_{t'}(w) + \beta(t)h(w) \right\} .$$

Proof. Since $z_i(1) = 0$ by definition, $\bar{z}(1) = 0$ also, and we have $\bar{z}(t) = \sum_{t'=1}^{t-1} \bar{g}(t')$. It follows that

$$\bar{w}(t) = \arg \min_{w \in \mathcal{W}} \left\{ \left\langle \sum_{t'=1}^{t-1} \bar{g}(t'), w \right\rangle + \beta(t)h(w) \right\} . \tag{24}$$

Now consider the inner product and observe that, by adding and subtracting $F(\bar{w}(t'))$ and $\hat{g}(t')$, we have

$$\left\langle \sum_{t'=1}^{t-1} \bar{g}(t'), w \right\rangle = \sum_{t'=1}^{t-1} [\langle \nabla_w F(\bar{w}(t')), w \rangle + \langle \hat{g}(t') - \nabla_w F(\bar{w}(t')), w \rangle + \langle \bar{g}(t') - \hat{g}(t'), w \rangle] .$$

Moreover, since $F(\bar{w}(t'))$ and $\bar{w}(t')$ do not depend on w , minimizing the objective in (24) is equivalent to minimizing

$$\sum_{t'=1}^{t-1} [F(\bar{w}(t')) + \langle \nabla_w F(\bar{w}(t')), w - \bar{w}(t') \rangle + \langle \hat{g}(t') - \nabla_w F(\bar{w}(t')), w - \bar{w}(t') \rangle + \langle \bar{g}(t') - \hat{g}(t'), w - \bar{w}(t') \rangle] + \beta(t)h(w)$$

over w , which is what we wanted to show. \square

The second lemma we will make use of is proved in [30] and states a widely-used property of strongly convex functions.

Lemma 6. Let \mathcal{W} be a closed convex set, let $\phi(w)$ be a convex function on \mathcal{W} , and let $\tilde{h}(w)$ be a μ -strongly convex function defined on \mathcal{W} . Define

$$w^+ \stackrel{\text{def}}{=} \arg \min_{w \in \mathcal{W}} \{\phi(w) + \tilde{h}(w)\}.$$

Then for all $w \in \mathcal{W}$,

$$\phi(w) + \tilde{h}(w) \geq \phi(w^+) + \tilde{h}(w^+) + \frac{\mu}{2} \|w - w^+\|^2.$$

Now consider Lemma 6 with $\phi(w) = \sum_{t'=1}^{t-1} \hat{l}_{t'}(w)$ and $\tilde{h}(w) = (K + a(t))h(w)$. Since $h(w)$ is 1-strongly convex, $\tilde{h}(w)$ is $(K + a(t))$ -strongly convex. Thus

$$\begin{aligned} \frac{K + a(t)}{2} \|\bar{w}(t+1) - \bar{w}(t)\|^2 &\leq \sum_{t'=1}^{t-1} \hat{l}_{t'}(\bar{w}(t+1)) + (K + a(t))h(\bar{w}(t+1)) \\ &\quad - \sum_{t'=1}^{t-1} \hat{l}_{t'}(\bar{w}(t)) - (K + a(t))h(\bar{w}(t)). \end{aligned}$$

Using this bound in (23) gives

$$\begin{aligned} F(\bar{w}(t+1)) &\leq \sum_{t'=1}^t \hat{l}_{t'}(\bar{w}(t+1)) + (K + a(t))h(\bar{w}(t+1)) \\ &\quad - \sum_{t'=1}^{t-1} \hat{l}_{t'}(\bar{w}(t)) - (K + a(t))h(\bar{w}(t)) \\ &\quad + \frac{\|q(t)\|^2}{4a(t)} + \frac{\|r(t)\|^2}{4a(t)}. \end{aligned}$$

Summing both sides over $t = 1, \dots, \frac{m}{b+\mu} - 1$ leads to many terms on the right-hand side cancelling, and we are left with

$$\begin{aligned} \sum_{t=2}^{\frac{m}{b+\mu}} F(\bar{w}(t)) &\leq \sum_{t=1}^{\frac{m}{b+\mu}-1} \hat{l}_t(\bar{w}(\frac{m}{b+\mu})) + (K + a(t))h(\bar{w}(\frac{m}{b+\mu})) \\ &\quad + \sum_{t=1}^{\frac{m}{b+\mu}-1} \frac{\|q(t)\|^2 + \|r(t)\|^2}{4a(t)}. \end{aligned}$$

From Lemma 5, it follows that

$$\sum_{t=1}^{\frac{m}{b+\mu}-1} \hat{l}_t(\bar{w}(\frac{m}{b+\mu})) + (K + a(t))h(\bar{w}(\frac{m}{b+\mu})) \leq \sum_{t=1}^{\frac{m}{b+\mu}-1} \hat{l}_t(w^*) + (K + a(\frac{m}{b+\mu}))h(w^*),$$

and so, replacing $\hat{l}_t(w)$ with the characterization given in (21), we have

$$\begin{aligned} \sum_{t=2}^{\frac{m}{b+\mu}} F(\bar{w}(t)) &\leq \sum_{t=1}^{\frac{m}{b+\mu}-1} l_t(w^*) + \sum_{t=1}^{\frac{m}{b+\mu}-1} [\langle q(t), w^* - \bar{w}(t) \rangle + \langle r(t), w^* - \bar{w}(t) \rangle] \\ &\quad + (K + a(\frac{m}{b+\mu}))h(w^*) + \sum_{t=1}^{\frac{m}{b+\mu}-1} \frac{\|q(t)\|^2 + \|r(t)\|^2}{4a(t)} \\ &\leq \left(\frac{m}{b+\mu} - 1 \right) F(w^*) + \sum_{t=1}^{\frac{m}{b+\mu}-1} [\langle q(t), w^* - \bar{w}(t) \rangle + \langle r(t), w^* - \bar{w}(t) \rangle] \\ &\quad + (K + a(\frac{m}{b+\mu}))h(w^*) + \sum_{t=1}^{\frac{m}{b+\mu}-1} \frac{\|q(t)\|^2 + \|r(t)\|^2}{4a(t)}. \end{aligned}$$

Add $F(\bar{w}(1)) - F(w^*)$ to both sides, take the expectation with respect to the data $X(\frac{m}{b+\mu})$, and apply the bounds given in Lemma 4 to obtain the desired bound,

$$\begin{aligned}
\sum_{t=1}^{\frac{m}{b+\mu}} \mathbb{E}[F(\bar{w}(t)) - F(w^*)] &\leq F(\bar{w}(1)) - F(w^*) + (K + a(\frac{m}{b+\mu}))h(w^*) + \sum_{t=1}^{\frac{m}{b+\mu}-1} \frac{\sigma^2/b}{4a(t)} \\
&\quad + \frac{K^2}{4n^2} \sum_{t=1}^{\frac{m}{b+\mu}-1} \frac{1}{a(t)(K + a(t))^2} \sum_{i=1}^n \sum_{j=1}^n \mathbb{E}[\|z_i(t) - \bar{z}(t)\| \cdot \|z_j(t) - \bar{z}(t)\|] \\
&\quad + \frac{KD}{n} \sum_{t=1}^{\frac{m}{b+\mu}-1} \frac{1}{K + a(t)} \mathbb{E}[\|z_i(t) - \bar{z}(t)\|] \\
&\leq F(\bar{w}(1)) - F(w^*) + (K + a(\frac{m}{b+\mu}))h(w^*) + \sum_{t=1}^{\frac{m}{b+\mu}-1} \frac{\sigma^2/b}{4a(t)} \\
&\quad + \frac{\delta^2 K^2}{4} \sum_{t=1}^{\frac{m}{b+\mu}-1} \frac{1}{a(t)(K + a(t))^2} \\
&\quad + \frac{\delta KD}{n} \sum_{t=1}^{\frac{m}{b+\mu}-1} \frac{1}{K + a(t)}.
\end{aligned}$$

6.3 Proof of Theorem 2

Recall from Lemma 1 that, given initial values $y_j^{(0)} \in \mathbb{R}^d$ at each node $j \in [n]$ and given a constant $\delta > 0$, the distance between the average $\bar{y} = \frac{1}{n} \sum_{j=1}^n y_j^{(0)}$ and the value $y_j^{(k)}$ at each node after k iterations satisfies

$$\max_j \|y_j^{(k)} - \bar{y}\| \leq \delta$$

if

$$k \geq \frac{\log\left(\frac{1}{\delta} \cdot 2\sqrt{n} \cdot \max_j \|y_j^{(0)} - \bar{y}\|\right)}{1 - \lambda_2(P)}. \quad (25)$$

We wish to determine a condition on the number of gossip iterations k such that the accuracy δ and delay μ satisfy $\delta \leq \frac{1}{b+\mu} = \frac{1}{b+\gamma k}$.

When gossip is used for distributed averaging within the algorithm (4)–(6) to compute the values $\{z_j(t+1)\}_{j \in [n]}$ at round $t+1$, we have initial value $y_j^{(0)} = z_j(t) + g_j(t)$ at node j , and thus $\bar{y} = \bar{z}(t) + \bar{g}(t)$, where

$$\bar{z}(t) = \frac{1}{n} \sum_{j=1}^n z_j(t) \quad \text{and} \quad \bar{g}(t) = \frac{1}{n} \sum_{j=1}^n g_j(t).$$

Note that $\|\bar{g}(t)\| \leq L$ since $\|g_j(t)\| \leq L$ for all $j \in [n]$ by Assumption 2.

To show that the accuracy achieved is $\delta \leq \frac{1}{b+\gamma k}$, we must show that $\|z_j(t) - \bar{z}(t)\| \leq \delta$ for all $j \in [n]$ and all t . We will show this inductively, while also deriving the required number of gossip iterations k . For the induction base, note that the algorithm is initialized with $z_j(1) = 0$ for all $j \in [n]$, and so $\|z_j(1) - \bar{z}(1)\| = 0$. For the induction step, assume that $\|z_j(t) - \bar{z}(t)\| \leq \delta$ for all $j \in [n]$. Then we have that

$$\begin{aligned}
\max_j \|y_j^{(0)} - \bar{y}\| &= \max_j \|(z_j(t) + g_j(t)) - (\bar{z}(t) + \bar{g}(t))\| \\
&\leq \max_j \{\|z_j(t) - \bar{z}(t)\| + \|g_j(t) - \bar{g}(t)\|\} \\
&\leq \max_j \{\delta + 2L\} \\
&\leq \frac{1}{b + \gamma k} + 2L.
\end{aligned}$$

Using this together with the relation $\frac{1}{\delta} = b + \gamma k$ in (25), our task becomes that of finding the smallest integer k such that

$$k \geq \frac{\log(2\sqrt{n}(1 + 2Lb + 2L\gamma k))}{1 - \lambda_2(P)}.$$

For $x \in \mathbb{R}$, $x > \frac{-(1+2Lb)}{2L\gamma}$, define the function

$$\phi(x) \stackrel{\text{def}}{=} \frac{\log(2\sqrt{n}(1 + 2Lb + 2L\gamma x))}{1 - \lambda_2(P)}.$$

Since $\frac{1}{1-\lambda_2(P)}$, δ , L , n , and b are all positive constants by assumption, it follows that $\phi(0) > 0$. Therefore, $\phi(x)$ intersects the diagonal $\psi(x) = x$ at exactly one point $x > 0$. The exact solution to the equation $x = \phi(x)$ is

$$x = \frac{-1}{1 - \lambda_2(P)} \cdot W\left(-\frac{(1 - \lambda_2(P))}{4L\gamma\sqrt{n}} \exp\left\{-\frac{(1 - \lambda_2(P))(1 + 2Lb)}{(2L\gamma)}\right\}\right) - \frac{1 + 2Lb}{2L\gamma},$$

where $W(\cdot)$ is the Lambert W-function.⁶ Given b , L , n , γ , and $\lambda_2(P)$, one could numerically determine an appropriate value for k , but because $W(\cdot)$ is transcendental, this is not directly useful for obtaining the closed-form expression we desire.

Instead, since $\phi(x)$ is concave in x , it holds that for any $\hat{x} > \frac{-(1+2Lb)}{2L\gamma}$,

$$\begin{aligned} \phi(x) &\leq \phi(\hat{x}) + \phi'(\hat{x}) \cdot (x - \hat{x}) \\ &= \frac{\log(2\sqrt{n}(1 + 2Lb + 2L\gamma\hat{x}))}{1 - \lambda_2(P)} + \frac{L\gamma(x - \hat{x})}{(1 - \lambda_2(P))(1 + 2Lb + 2L\gamma\hat{x})}. \end{aligned}$$

Thus, $x \geq \phi(x)$ if $x \geq x^*$ where

$$x^* = \left(\frac{\log(2\sqrt{n}(1 + 2Lb + 2L\gamma\hat{x}))}{1 - \lambda_2(P)} - \frac{2L\gamma\hat{x}}{(1 - \lambda_2(P))(1 + 2Lb + 2L\gamma\hat{x})} \right) \cdot \frac{1}{1 - \frac{2L\gamma}{(1 - \lambda_2(P))(1 + 2Lb + 2L\gamma\hat{x})}}.$$

Now, take

$$\hat{x} = \frac{1}{2L\gamma} \left(\frac{2Lb}{1 - \lambda_2(P)} - 1 - 2Lb \right).$$

Then

$$\begin{aligned} x^* &= \left(\frac{\log(4\sqrt{n}Lb)}{1 - \lambda_2(P)} + \frac{\log(\frac{1}{1 - \lambda_2(P)})}{1 - \lambda_2(P)} - \frac{1}{1 - \lambda_2(P)} + \frac{1}{2Lb} + 1 \right) \frac{1}{1 - \frac{\gamma}{b}} \\ &\leq \left(\frac{\log(4Lb\sqrt{n})}{1 - \lambda_2(P)} + \frac{\log(\frac{1}{1 - \lambda_2(P)})}{1 - \lambda_2(P)} + \frac{1}{2Lb} + 1 \right) \frac{1}{1 - \frac{\gamma}{b}}. \end{aligned}$$

Rounding up to the next largest integer gives the value of k in the statement of the theorem. Thus, we have shown that $k^* = \lceil x^* \rceil$ satisfies the relation $k^* \geq \phi(k^*)$, and so

$$\max_j \|z_j(t+1) - \bar{z}(t+1)\| \leq \frac{1}{b + \gamma k} \leq \frac{1}{b + \mu}$$

holds, completing the induction step. It follows from Corollary 1 that the expected regret is bounded by $\mathbb{E}[R_n(m)] = \mathcal{O}(b + \sqrt{m})$, and so choosing $b = m^\rho$ for some $\rho \in (0, 1/2)$ gives that $\mathbb{E}[R_n(m)] = \mathcal{O}(\sqrt{m})$.

⁶The function $W(\cdot)$ is defined as the solution to the equation $W(x)e^{W(x)} = x$, and it has a real solution for $x \geq -1/e$; see [7] for more.

6.4 Proof of Theorem 3

Similar to the proof of Theorem 1, we will bound the optimization error uniformly over all nodes by comparing the performance of the sequence $\{\hat{w}_i(t)\}$ to that of an appropriate averaged sequence. Specifically, define

$$\widehat{\bar{w}}(t) \stackrel{\text{def}}{=} \frac{1}{t} \sum_{t'=1}^t \bar{w}(t'), \quad (26)$$

where $\bar{w}(t')$ is as defined in (15).

Recall that the network collectively processes b samples in each round of the algorithm, so that a total of $m = bT$ samples have been processed after T rounds. We seek an upper bound on $\mathbb{E}[F(\hat{w}_i(\frac{m}{b})) - F(w^*)]$. Adding and subtracting $F(\widehat{\bar{w}}(\frac{m}{b}))$, applying Jensen's inequality, and using Lipschitz continuity of $F(\cdot)$ gives

$$\begin{aligned} \mathbb{E}[F(\hat{w}_i(\frac{m}{b})) - F(w^*)] &= \mathbb{E}[F(\widehat{\bar{w}}(\frac{m}{b})) - F(w^*) + F(\hat{w}_i(\frac{m}{b})) - F(\widehat{\bar{w}}(\frac{m}{b}))] \\ &\leq \mathbb{E} \left[\frac{b}{m} \sum_{t=1}^{\frac{m}{b}} [F(\bar{w}(t)) - F(w^*)] + \frac{b}{m} \sum_{t=1}^{\frac{m}{b}} [F(w_i(t)) - F(\bar{w}(t))] \right] \\ &\leq \mathbb{E} \left[\frac{b}{m} \sum_{t=1}^{\frac{m}{b}} [F(\bar{w}(t)) - F(w^*)] + \frac{bL}{m} \sum_{t=1}^{\frac{m}{b}} \|w_i(t) - \bar{w}(t)\| \right]. \end{aligned} \quad (27)$$

Applying Lemma 3, we have that

$$\begin{aligned} \mathbb{E} \left[\frac{b}{m} \sum_{t=1}^{\frac{m}{b}} [F(\bar{w}(t)) - F(w^*)] \right] &\leq \frac{b}{m} \left[F(\bar{w}(1)) - F(w^*) + (K + a(\frac{m}{b}))h(w^*) + \frac{\sigma^2}{4b} \sum_{t=1}^{\frac{m}{b}} \frac{1}{a(t)} \right. \\ &\quad \left. + \frac{\delta^2 K^2}{4} \sum_{t=1}^{\frac{m}{b}} \frac{1}{a(t)(K + a(t))^2} + \frac{\delta KD}{n} \sum_{t=1}^{\frac{m}{b}} \frac{1}{K + a(t)} \right]. \end{aligned} \quad (28)$$

Also, using Lipschitz continuity of the proximal projection operation (cf. Lemma 2), we have that

$$\mathbb{E} \left[\frac{bL}{m} \sum_{t=1}^{\frac{m}{b}} \|w_i(t) - \bar{w}(t)\| \right] \leq \frac{bL}{m} \sum_{t=1}^{\frac{m}{b}} \frac{1}{\beta(t)} \|z_i(t) - \bar{z}(t)\|. \quad (29)$$

Combining (27), (28), and (29), and recalling that it is assumed in the statement of the theorem that the approximate distributed averaging protocol has accuracy $\|z_i(t) - \bar{z}(t)\| \leq \delta \leq \frac{1}{b}$, we obtain

$$\mathbb{E}[F(\hat{w}_i(\frac{m}{b})) - F(w^*)] \leq \frac{1}{T} [F(\bar{w}(1)) - F(w^*)] + \frac{1}{T} Kh(w^*) + \frac{\sigma^2}{4\sqrt{bT}} + \frac{3K^2}{4T\sqrt{b}} + \frac{2KD}{n\sqrt{bT}} + \frac{2L}{\sqrt{bT}}.$$

Using the fact that $b = Cn$, and rearranging terms gives the desired result.

7 Experiments

In this section we illustrate the performance of the proposed distributed dual averaging method through numerical experiments. The experiments are conducted on a cluster consisting of eight servers. Each server has two Intel Xeon 2.5GHz quad-core processors and 14GB of RAM. Our experiments make use of up to $n = 64$ processes, with each process (i.e., node in the terminology of the previous sections) running on a separate core. The servers communicate over 1Gbps Ethernet connections.

We consider a multi-class classification task. In this task, each data point is a input-output pair, $(x, y) \in \mathcal{X} \times \mathcal{Y}$, with the input x drawn from input space $\mathcal{X} \subseteq \mathbb{R}^d$, and the output y drawn from a discrete output space $\mathcal{Y} = \{1, 2, \dots, M\}$ denoting the class of the data point. We use the MNIST digits dataset. Each input

x is a 28×28 greyscale image corresponding to one of the digits $0, 1, \dots, 9$. Hence, the input space has dimension $d = 784$ and there are $M = 10$ classes.

We use multinomial logistic regression to address the multi-class classification task. In particular, the primal parameter w is $M(d + 1)$ dimensional and is broken into M blocks, $w = [w_1^T, w_2^T, \dots, w_M^T]^T$, with $w_c \in \mathbb{R}^{d+1}$ being the parameters for the c th class. The loss function $f(w, (x, y))$ is defined as

$$f(w, (x, y)) = \frac{1}{Z(w, x)} \exp \left\{ w_{y, d+1} + \sum_{j=1}^d w_{y, j} x_j \right\},$$

where the normalization constant is given by

$$Z(w, x) = \sum_{k=1}^M \exp \left\{ w_{k, d+1} + \sum_{j=1}^d w_{k, j} x_j \right\}.$$

We report two sets of experiments, one in which we keep the mini-batch size b fixed and vary the number of nodes, and the other where we fix the ratio b/n (i.e., the number of data points processed by each node per round) and vary the number of nodes. In both experiments, gossip-based distributed dual averaging is implemented using the Matlab Parallel Computing Toolbox and Distributed Computing Server. The communication topology G is obtained by sampling a random graph from the Erdős-Rényi model with $p = 0.5$; all nodes have, on average, the same number of neighbours and the graph is approximately a constant-degree expander.

7.1 Fixed mini-batch size

In this experiment we fix the mini-batch size b to be equal to the total number of data points available (roughly 50k) and vary the number of nodes. Theorem 2 states that the regret after $T = \frac{m}{b+\mu}$ rounds should scale like $\mathbb{E}[R_n(m)] = \mathcal{O}(\sqrt{m})$. At each round, each node processes b/n data points, so we have that $m = T \cdot \frac{b}{n} \cdot n = Tb$. Hence, the expected regret per data point should decay as $\frac{1}{m} \mathbb{E}[R_n(m)] = \mathcal{O}(\frac{1}{\sqrt{Tb}})$; i.e., this error only depends on the number of data points observed and not on the size of the network. We have observed exactly this in our experiments. Varying $n = 2^k$ for $k = 2, 3, \dots, 6$, and plotting $\frac{1}{m} \mathbb{E}[R_n(T)]$ versus T , the resulting curves all lie directly on top of each other; the figure is omitted since it provides no additional insight.

The benefit of using more nodes in this setting is that the workload of computing the batch gradient over b data points is distributed over multiple processors, and so each round should require less time. Figure 1 shows $\frac{1}{m} \mathbb{E}[R_n(m)]$ as a function of run time for varying number of processors. Increasing the number of processors gives faster convergence to a lower average risk per data point, as expected. However, the benefit of adding more processors seems to obey a law of diminishing returns. This is related to the fact that using more processors requires more communication, and communication also requires time [31].

In this experiment, the network only performs one consensus iteration per round, rather than the number prescribed by Theorem 2. In other experiments, not reported here, we have observed that using additional consensus iterations does not significantly change the results.

7.2 Varying mini-batch size

In the next set of experiments, we set the mini-batch size to be $b = 200 \times n$ (i.e., each node processes 200 data points per round), and we again vary the number of nodes as $n = 2^k$ for $k = 2, 3, \dots, 6$. In this case, since each node processes the same amount of data per round, the time per round should be roughly constant, regardless of the network size. However, after a fixed number of rounds, T , the regret per data point $\frac{1}{m} \mathbb{E}[R_n(T)]$ should be lower for larger networks since they will have seen more cumulative data. In particular, with b/n fixed constant, we have $\mathbb{E}[R_n(T)] = \mathcal{O}(\sqrt{Tn})$, and so for $n' > n$, we expect that the ratio $R_n(T)/R_{n'}(T)$ should roughly behave as $\mathcal{O}(\sqrt{n/n'})$. The figure clearly shows that this behavior arises for sufficiently large T , after the transients have vanished and the $\mathcal{O}(\sqrt{m})$ behavior dominates.




Figure 1: Regret per data point as a function of time, for varying number of processors with fixed batch size $b = 49984$.




Figure 2: Ratio of regrets $R_4(T)/R_n(T)$ versus number of rounds T for $n = 8, 16, 32$ with batch size $b = 200n$ proportional to the network size. Note, $\sqrt{2} \approx 1.4$, $\sqrt{4} = 2$, and $\sqrt{8} \approx 2.8$.

In these experiments, similar to above, only a single consensus iteration is performed over the network per round. In the proof of Theorem 2, the number of consensus iterations given in (10) is required to ensure that the dual variables at all nodes are sufficiently close to the average of the dual variables, $\frac{1}{n} \sum_{i=1}^n z_i(t)$. In particular, this is the worst-case number of iterations required for any initial spread of the $z_i(t)$. However, in a typical execution of the algorithm, after processing one mini-batch of data the values at each node will not vary too greatly, and so the initial disagreement is not so large. Hence, we conjecture that an adaptive algorithm, which guarantees that the disagreement (i.e., deviation from the average) at each iteration is not too large, should be sufficient to obtain scaling results comparable to those in Theorem 2.

8 Discussion

In comparison to the distributed mini-batch algorithm [8], we have shown that the proposed distributed dual averaging algorithm with approximate mini-batches achieves the same order-wise performance. In particular, if the communication topology is an expander graph and the mini-batch size is chosen to be $b = \Theta(\sqrt{m})$, then the expected regret is bounded as $\mathbb{E}[R_n(m)] = \mathcal{O}(\sqrt{m})$, and the time spent communicating in each round is of the order $\mathcal{O}(\log n)$. This is comparable to the time spent performing exact distributed averaging using a protocol such as ALLREDUCE, but gossip protocols can be considerably more robust to node and link failures [14, 32].

Raginsky et al. [23] and Yan et al. [37] study the related problem of distributed online convex optimization, where a different (deterministic) cost function f_t is chosen by an adversary at each round, and the goal of the network is to minimize the regret with respect to the decision variable w that would have minimized the overall cost $\sum_t f_t(w)$ in hindsight. The approach of Raginsky et al. [23] also builds on the dual averaging algorithm; however, rather than fusing information through gossip iterations, their approach involves the solution of a local dynamic program at each node.

Method	Num. Rounds (T)	Gossip Iter. / Round	Total Gossip Iter.
[17, 24]	$\mathcal{O}\left(\frac{n^3}{\epsilon^2}\right)$	1	$\mathcal{O}\left(\frac{n^3}{\epsilon^2}\right)$
[10]	$\mathcal{O}\left(\frac{1}{\epsilon^2} \frac{\log n}{\sqrt{1-\lambda_2(P)}}\right)$	1	$\mathcal{O}\left(\frac{1}{\epsilon^2} \frac{\log n}{\sqrt{1-\lambda_2(P)}}\right)$
This Paper	$\mathcal{O}\left(\frac{1}{n\epsilon^2}\right)$	$\mathcal{O}\left(\frac{\log n}{1-\lambda_2(P)}\right)$	$\mathcal{O}\left(\frac{1}{\epsilon^2} \frac{\log n}{n(1-\lambda_2(P))}\right)$

Table 1: Comparing the number of rounds and gossip iterations required to guarantee an optimality gap of ϵ for distributed stochastic optimization.

This paper also makes contributions to the literature on gossip-based distributed optimization algorithms. The use of distributed averaging algorithms as a means of coordination in distributed optimization methods has a history tracing back to the seminal work of Tsitsiklis et al. [34]. Distributed subgradient methods [17, 24] that use gossip protocols to aggregate primal variables generally come with bounds guaranteeing that the optimality gap decays as $\mathcal{O}\left(n^{3/2}/\sqrt{T}\right)$. The poor scaling in the network size appears since the analysis for these methods is effectively over the worst-case network topology. The distributed dual averaging algorithm described in [10] guarantees that the error is less than $\mathcal{O}\left(\frac{\log(T\sqrt{n})}{\sqrt{T \cdot (1-\lambda_2(P))}}\right)$ after T iterations. The number of transmissions per node is proportional to the number of iterations for the algorithms [10, 17, 24]. In contrast, for the proposed approach, the error decays as $\mathcal{O}\left(\frac{1}{\sqrt{nT}}\right)$ where T is the number of rounds, and in each round the network performs $\mathcal{O}\left(\frac{\log(n)}{1-\lambda_2(P)}\right)$ gossip iterations. Consequently, the overall communications required to reach accuracy ϵ is reduced by a factor of $\mathcal{O}(n)$, as illustrated in Table 1. This is not surprising, since our analysis is based on stronger assumptions than [10, 17, 24]; in particular, we assume that the loss function has Lipschitz continuous gradients, whereas the previous work only studied loss functions which are convex, Lipschitz continuous, and have bounded gradients. The proposed method has a number of gossip iterations per round that depends on the spectral gap via $\frac{1}{1-\lambda_2(P)}$. This can be immediately reduced to $\frac{1}{\sqrt{1-\lambda_2(P)}}$ by using the accelerated distributed averaging method of [22], while only requiring that each node stores one extra copy of the dual variable being averaged.

As a side remark, we point out that faster gossip-based algorithms exist for more restricted function classes. Recall that in the setting of stochastic optimization the constant multiplying the $\mathcal{O}\left(\frac{1}{\sqrt{T}}\right)$ term in single-processor optimality gap bounds is proportional to the standard deviation of the stochastic gradient magnitudes [16, 35]. In the deterministic, single-processor setting, and under the assumption that the loss function has Lipschitz continuous gradients, it is known that significantly faster rates are possible [18]. A gossip-based algorithm which achieves rates of $G_n(T) = \mathcal{O}\left(\frac{1}{T^2}\right)$ is described in [13]. Similar to the gossip-based approach proposed in this paper, the approach in [13] also requires performing multiple gossip iterations per round. Unlike the approach in this paper, however, the number of gossip iterations per round increases as more rounds are performed; specifically, the number of gossip iterations required in round t is $\mathcal{O}\left(\frac{\log(t)}{1-\lambda_2(P)}\right)$. In addition, the method of [13] must know the values of the gradient Lipschitz constant K and the spectral gap $1 - \lambda_2(P)$ in order to determine a step-size sequence that achieves these bounds.

Finally, we note that similar rates for distributed stochastic optimization are achieved by a method based on a master-worker architecture [1]. In this architecture, the master maintains and updates the authoritative copy of the optimization variables. The workers compute mini-batch gradients and report these back to the master, where they are used to update the optimization variables. The challenge in this setup is that the workers report gradients evaluated at out-dated values of $w(t)$. The analysis in [1] involves bounding and controlling the additional error due to these delayed evaluations, and the resulting bounds are, at best, of the form $\Delta_n(T) = \mathcal{O}\left(\frac{n}{T} + \frac{1}{\sqrt{nT}}\right)$ or $\mathcal{O}\left(\frac{1}{T^{2/3}} + \frac{1}{\sqrt{nT}}\right)$, depending on the version of the algorithm considered. There are also potential practical issues which must be addressed when implementing a master-worker architecture; for example, the master may become a bottleneck and single point of failure.

9 Conclusion and Future Work

The main message of this paper is that exact distributed averaging is not essential to achieve optimal rates with mini-batch methods for distributed online prediction and stochastic optimization. This opens the door to the use of gossip protocols for distributed averaging, which can be considerably more robust than exact distributed averaging protocols and can also be implemented in a completely asynchronous manner. A natural extension, and an important part of our ongoing work, is to analyze a completely asynchronous version of the distributed dual averaging algorithm with approximate mini-batch calculations. In the asynchronous setting, nodes may send and receive a random number of gossip messages, and they may also update their local copies of the decision variables at different rates. There are two major challenges in this setting. The first is to ensure that performance is not seriously degraded if, occasionally, an insufficient amount of gossiping occurs between subsequent updates. The second is to properly adapt the algorithm parameters $\beta(t)$ in such a manner that the asynchronous algorithm still converges at the optimal rate.

The algorithm and analysis in this paper are appropriate when the loss function is convex and has Lipschitz continuous gradients. For that class of problems, the proposed algorithm achieves the best rates possible. However, for other classes of loss functions — in particular, for strongly convex loss functions — it is known that better rates (e.g., $\mathbb{E}[R_1(m)] = \mathcal{O}(\log m)$) are achievable in the single-processor setting [2, 29, 30]. At present, there are no known distributed optimization methods that achieve comparable rates while also exhibiting scalability in terms of the network size.

A Proof of Lemma 1

The proof of Lemma 1 makes use of the following fact for stochastic matrices; see, e.g., [12].

Lemma 7. *Let P be an $n \times n$ doubly-stochastic matrix with eigenvalues $\lambda_1(P) = 1 \geq \lambda_2(P) \geq \dots \geq \lambda_n(P) \geq 0$, and let z be a vector in the n -dimensional unit simplex. Then*

$$\left\| z^T P^k - \frac{1}{n} \mathbf{1}_n^T \right\|_1 \leq \lambda_2(P)^k \sqrt{n}.$$

Recall that $\bar{y} = \frac{1}{n} \sum_{j=1}^n y_j^{(0)}$. After k synchronous iterations of gossip, for all $i \in [n]$ it holds that

$$\begin{aligned} y_i^{(k)} - \bar{y} &= \sum_{j=1}^n [P^k]_{i,j} (y_j^{(0)} - \bar{y}) \\ &= \sum_{j=1}^n \left([P^k]_{i,j} - \frac{1}{n} \right) y_j^{(0)} \\ &= \sum_{j=1}^n \left([P^k]_{i,j} - \frac{1}{n} \right) (y_j^{(0)} - \bar{y}), \end{aligned}$$

where the last line follows since P^k is also doubly stochastic, and so $\sum_{j=1}^n [P^k]_{i,j} = \sum_{i=1}^n \frac{1}{n} = 1$.

Applying the triangle inequality gives

$$\begin{aligned} \|y_i^{(k)} - \bar{y}\| &\leq \sum_{j=1}^n \left| [P^k]_{i,j} - \frac{1}{n} \right| \cdot \|y_j^{(0)} - \bar{y}\| \\ &\leq \left(\max_j \|y_j^{(0)} - \bar{y}\| \right) \left\| \mathbf{e}_i^T P^k - \frac{1}{n} \mathbf{1}_n^T \right\|_1, \end{aligned}$$

where \mathbf{e}_i denotes the i th canonical vector, and where $\|\cdot\|_1$ denotes the ℓ_1 norm; $\|x\|_1 = \sum_i |x_i|$.⁷ Clearly \mathbf{e}_i lies in the unit simplex. Applying the result from Lemma 7, and noting that $\lambda_2(P^k) = \lambda_2(P)^k$, we have that

$$\|y_i^{(k)} - \bar{y}\| \leq 2\sqrt{n} \lambda_2(P)^k \cdot \max_j \|y_j^{(0)} - \bar{y}\|.$$

⁷I.e., $\mathbf{e}_i \in \{0, 1\}^n$ has j th entry $[\mathbf{e}_i]_j = 1$ if $j = i$, and $[\mathbf{e}_i]_j = 0$ otherwise.

It follows that $\left\| y_i^{(k)} - \bar{y} \right\| \leq \delta$ if

$$k \geq \frac{\log \left(\frac{1}{\delta} \cdot 2\sqrt{n} \cdot \max_j \left\| y_j^{(0)} - \bar{y} \right\| \right)}{\log(1/\lambda_2(P))}.$$

The claim of the lemma follows by recalling that $\log(1/x) \geq 1 - x$ for $x > 0$.

B Proof of Lemma 4

We prove the four relationships in the order they are stated in the lemma. First consider $\mathbb{E}[\langle q(t), w^* - \bar{w}(t) \rangle]$. Notice that $\bar{w}(t)$ is a deterministic function given $X(t-1)$ so

$$\mathbb{E}[f(\bar{w}(t), x_i(t))] = \mathbb{E}[\mathbb{E}[f(\bar{w}(t), x_i(t)) | X(t-1)]] = \mathbb{E}[F(\bar{w}(t))].$$

Also, under Assumption 2, the gradient and expectation operators commute [26], and so we have

$$\mathbb{E}[\langle \nabla f(\bar{w}(t), x_j(t)), \bar{w}(t) \rangle] = \mathbb{E}[\langle \nabla F(\bar{w}(t)), \bar{w}(t) \rangle].$$

Therefore, for the first expression in the lemma, we get

$$\begin{aligned} \mathbb{E}[\langle q(t), w^* - \bar{w}(t) \rangle] &= \mathbb{E}[\langle \hat{g}(t) - \nabla F(\bar{w}(t)), w^* - \bar{w}(t) \rangle] \\ &= \mathbb{E} \left[\left\langle \frac{1}{b} \sum_{i=1}^n \sum_{s=1}^{\frac{b}{n}} \nabla f(\bar{w}(t), x_i(t, s)) - \nabla F(\bar{w}(t)), w^* - \bar{w}(t) \right\rangle \right] \\ &= \mathbb{E} \left[\left\langle \frac{1}{b} \sum_{i=1}^n \sum_{s=1}^{\frac{b}{n}} \nabla F(\bar{w}(t)) - \nabla F(\bar{w}(t)), w^* - \bar{w}(t) \right\rangle \right] \\ &= 0. \end{aligned}$$

For the second expression, using similar conditioning arguments gives

$$\begin{aligned} \mathbb{E}[\langle r(t), w^* - \bar{w}(t) \rangle] &= \mathbb{E} \left[\left\langle \frac{1}{b} \sum_{i=1}^n \sum_{s=1}^{\frac{b}{n}} \nabla f(w_i(t), x_i(t, s)) - \frac{1}{b} \sum_{i=1}^n \sum_{s=1}^{\frac{b}{n}} \nabla f(\bar{w}(t), x_i(t, s)), w^* - \bar{w}(t) \right\rangle \right] \\ &= \mathbb{E} \left[\left\langle \frac{1}{b} \sum_{i=1}^n \sum_{s=1}^{\frac{b}{n}} \nabla F(w_i(t)) - \frac{1}{b} \sum_{i=1}^n \sum_{s=1}^{\frac{b}{n}} \nabla F(\bar{w}(t)), w^* - \bar{w}(t) \right\rangle \right] \\ &= \frac{1}{n} \sum_{i=1}^n \mathbb{E}[\langle \nabla F(w_i(t)) - \nabla F(\bar{w}(t)), w^* - \bar{w}(t) \rangle]. \end{aligned}$$

Then, applying the Cauchy-Schwarz inequality and using Assumptions 1 and 2 gives

$$\begin{aligned} \mathbb{E}[\langle r(t), w^* - \bar{w}(t) \rangle] &\leq \frac{1}{n} \sum_{i=1}^n \mathbb{E}[\|\nabla F(w_i(t)) - \nabla F(\bar{w}(t))\| \cdot \|w^* - \bar{w}(t)\|] \\ &\leq \frac{1}{n} \sum_{i=1}^n \mathbb{E}[K \|w_i(t) - \bar{w}(t)\| \cdot D]. \end{aligned}$$

Finally, applying Lemma 2 gives the desired result,

$$\mathbb{E}[\langle r(t), w^* - \bar{w}(t) \rangle] \leq \frac{KD}{n\beta(t)} \sum_{i=1}^n \mathbb{E}[\|z_i(t) - \bar{z}(t)\|].$$

The third expression follows by again using similar conditioning arguments and by applying Assumption 3. In particular, recalling that the data are i.i.d., and hence $x_i(t, s)$ is independent of $x_j(t, s)$, we have

$$\begin{aligned}
\mathbb{E} [\|q(t)\|^2] &= \mathbb{E} [\|\hat{g}(t) - \nabla F(\bar{w}(t))\|^2] \\
&= \mathbb{E} \left[\left\| \frac{1}{b} \sum_{i=1}^n \sum_{s=1}^{\frac{b}{n}} \nabla f(\bar{w}(t), x_i(t, s)) - \nabla F(\bar{w}(t)) \right\|^2 \right] \\
&= \mathbb{E} \left[\left\| \frac{1}{n} \sum_{i=1}^n \frac{n}{b} \sum_{s=1}^{\frac{b}{n}} \nabla f(\bar{w}(t), x_i(t, s)) - \nabla F(\bar{w}(t)) \right\|^2 \right] \\
&= \mathbb{E} \left[\frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \left\langle \frac{n}{b} \sum_{s=1}^{\frac{b}{n}} \nabla f(\bar{w}(t), x_i(t, s)) - \nabla F(\bar{w}(t)), \frac{n}{b} \sum_{s=1}^{\frac{b}{n}} \nabla f(\bar{w}(t), x_j(t, s)) - \nabla F(\bar{w}(t)) \right\rangle \right] \\
&= \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \mathbb{E} \left[\left\langle \frac{n}{b} \sum_{s=1}^{\frac{b}{n}} \nabla f(\bar{w}(t), x_i(t, s)) - \nabla F(\bar{w}(t)), \frac{n}{b} \sum_{s=1}^{\frac{b}{n}} \nabla f(\bar{w}(t), x_j(t, s)) - \nabla F(\bar{w}(t)) \right\rangle \right] \\
&= \frac{1}{n^2} \sum_{i=1}^n \mathbb{E} \left[\left\| \frac{n}{b} \sum_{s=1}^{\frac{b}{n}} \nabla f(\bar{w}(t), x_i(t, s)) - \nabla F(\bar{w}(t)) \right\|^2 \right].
\end{aligned}$$

Then, applying the triangle inequality and using Assumption 3 gives

$$\begin{aligned}
\mathbb{E} [\|q(t)\|^2] &\leq \frac{1}{n^2} \sum_{i=1}^n \frac{n^2}{b^2} \sum_{s=1}^{\frac{b}{n}} \mathbb{E} [\|\nabla f(\bar{w}(t), x_i(t, s)) - \nabla F(\bar{w}(t))\|^2] \\
&\leq \frac{1}{b^2} \sum_{i=1}^n \sum_{s=1}^{\frac{b}{n}} \sigma^2 \\
&\leq \frac{\sigma^2}{b}.
\end{aligned}$$

Finally, to verify the fourth expression, we first multiply and divide by n and use the triangle inequality to see that

$$\begin{aligned}
\mathbb{E} [\|r(t)\|^2] &= \mathbb{E} \left[\left\| \frac{1}{b} \sum_{i=1}^n \sum_{s=1}^{\frac{b}{n}} \nabla f(w_i(t), x_i(t, s)) - \nabla f(\bar{w}(t), x_i(t, s)) \right\|^2 \right] \\
&= \mathbb{E} \left[\left\| \frac{1}{n} \sum_{i=1}^n \frac{n}{b} \sum_{s=1}^{\frac{b}{n}} \nabla f(w_i(t), x_i(t, s)) - \nabla f(\bar{w}(t), x_i(t, s)) \right\|^2 \right] \\
&\leq \mathbb{E} \left[\left(\frac{1}{n} \sum_{i=1}^n \frac{n}{b} \sum_{s=1}^{\frac{b}{n}} \|\nabla f(w_i(t), x_i(t, s)) - \nabla f(\bar{w}(t), x_i(t, s))\| \right)^2 \right].
\end{aligned}$$

Then applying Assumption 2 and Lemma 2 and simplifying gives the desired result,

$$\begin{aligned}
\mathbb{E} [\|r(t)\|^2] &\leq \mathbb{E} \left[\left(\frac{1}{n} \sum_{i=1}^n \frac{n}{b} \sum_{s=1}^{\frac{b}{n}} K \|w_i(t) - \bar{w}(t)\| \right)^2 \right] \\
&\leq \mathbb{E} \left[\left(\frac{1}{n} \sum_{i=1}^n \frac{K}{\beta(t)} \|z_i(t) - \bar{z}(t)\| \right)^2 \right] \\
&= \mathbb{E} \left[\left(\frac{K}{n\beta(t)} \sum_{i=1}^n \|z_i(t) - \bar{z}(t)\| \right)^2 \right] \\
&= \mathbb{E} \left[\frac{K^2}{n^2\beta(t)^2} \sum_{i=1}^n \sum_{j=1}^n \|z_i(t) - \bar{z}(t)\| \|z_j(t) - \bar{z}(t)\| \right].
\end{aligned}$$

References

- [1] A. Agarwal and J. C. Duchi. Distributed delayed stochastic optimization. In *Neural Information Processing Systems*, 2011.
- [2] L. Bottou. Large-scale machine learning with stochastic gradient descent. In Y. Lechevallier and G. Saporta, editors, *Proceedings of the 19th International Conference on Computational Statistics*, pages 177–187, Paris, France, August 2010.
- [3] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah. Randomized gossip algorithms. *IEEE Transactions on Information Theory*, 52:2508–2530, 2006.
- [4] N. Cesa-Bianchi and G. Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006.
- [5] J. Chen and A. H. Sayed. Diffusion adaptation strategies for distributed optimization and learning over networks. *IEEE Transactions on Signal Processing*, 60(8):4289–4305, August 2012.
- [6] F. Chung. *Spectral Graph Theory*. AMS, 1998.
- [7] R. M. Corless, G. H. Gonnet, D. E. Hare, D. J. Jeffrey, and D. E. Knuth. On the Lambert W function. *Advances in Computational Mathematics*, 5(1):329–359, 1996.
- [8] O. Dekel, R. Gilad-Bachrach, O. Shamir, and L. Xiao. Optimal distributed online prediction using mini-batches. *Journal of Machine Learning Research*, 13:165–202, 2012.
- [9] A. G. Dimakis, S. Kar, J. M. Moura, M. G. Rabbat, and A. Scaglione. Gossip algorithms for distributed signal processing. *Proceedings of the IEEE*, 98(11):1847–1864, November 2010.
- [10] J. Duchi, A. Agarwal, and M. Wainwright. Dual averaging for distributed optimization: Convergence analysis and network scaling. *IEEE Transactions on Automatic Control*, 57(3):592–606, 2011.
- [11] W. Gropp, S. Huss-Lederman, A. Lumsdaine, E. Lusk, B. Nitzberg, W. Saphir, and M. Snir. *MPI—The Complete Reference - Volumes 1,2*. MIT Press, Cambridge, MA, 1998.
- [12] R. Horn and C. Johnson. *Matrix Analysis*. Cambridge University Press, 1990.
- [13] D. Jakovetic, J. Xavier, and J. M. Moura. Fast distributed gradient methods. *arXiv:1112.2972v3*, 2013.
- [14] M. Jelasity, A. Montessoro, and O. Babaoglu. Gossip-based aggregation in large dynamic networks. *ACM Transactions on Computer Systems*, 23(3):219–252, 2005.
- [15] D. Kempe, A. Dobra, and J. Gehrke. Gossip-based computation of aggregate information. In *FOCS*, vol. 44. IEEE Computer Society Press, pp. 482–491, 2003.

- [16] G. Lan. An optimal method for stochastic composite optimization. *Mathematical Programming*, 133(1):365–397, 2012.
- [17] A. Nedic and A. Ozdaglar. Distributed subgradient methods for multi-agent optimization. *IEEE Transactions on Automatic Control*, 54(1), January 2009.
- [18] A. Nemirovsky and D. Yudin. *Problem Complexity and Method Efficiency in Optimization*. John Wiley & Sons, 1983.
- [19] Y. Nesterov. *Introductory Lectures on Convex Optimization: A Basic Course*. Kluwer, Boston, 2004.
- [20] Y. Nesterov. Primal-dual subgradient methods for convex problems. *Mathematical Programming Series B*, 120:221–259, 2009.
- [21] R. Olfati-Saber, J. A. Fax, and R. M. Murray. Consensus and cooperation in networked multi-agent systems. In *Proceedings of the IEEE*, volume 95:1, pages 215 – 233, 2007.
- [22] B. Oreshkin, M. Coates, and M. Rabbat. Optimization and analysis of distributed averaging with short node memory. *IEEE Transactions on Signal Processing*, 58(5):2850–2865, May 2010.
- [23] M. Raginsky, N. Kiarashi, and R. Willett. Decentralized online convex programming with local information. In *Proceedings of the American Control Conference*, San Francisco, USA, 2011.
- [24] S. S. Ram, A. Nedic, and V. V. Veeravalli. Distributed stochastic subgradient projection algorithms for convex optimization. *Journal of Optimization Theory and Applications*, 147(3):516–545, 2011.
- [25] O. Reingold, S. Vadhan, and A. Wigderson. Entropy waves, the zig-zag graph product, and new constant-degree expanders. *Annals of Mathematics*, 155(2):157–187, 2002.
- [26] R. T. Rockafellar and R. J.-B. Wets. On the interchange of subdifferentiation and conditional expectation for convex functionals. *Stochastics: An International Journal of Probability and Stochastic Processes*, 7(3):173–182, 1982.
- [27] E. Seneta. *Non-negative Matrices and Markov Chains*. Springer, 1973.
- [28] S. Shalev-Shwartz. *Online learning and online convex optimization*. Foundations and Trends in Machine Learning, 4, (2):107-194, 2012.
- [29] S. Shalev-Shwartz and Y. Singer. Logarithmic regret algorithms for strongly convex repeated games. Technical report, The Hebrew University, 2007.
- [30] P. Tseng. On accelerated proximal gradient methods for convex-concave optimization. *SIAM Journal on Optimization*, 2008.
- [31] K. I. Tsianos, S. Lawlor, and M. G. Rabbat. Communication/computation tradeoffs in consensus-based distributed optimization. In *Neural Information Processing Systems*, 2012.
- [32] K. I. Tsianos, S. Lawlor, and M. G. Rabbat. Consensus-based distributed optimization: Practical issues and applications in large-scale machine learning. In *50th Allerton Conference on Communication, Control, and Computing*, 2012.
- [33] K. I. Tsianos, S. Lawlor, and M. G. Rabbat. Push-sum distributed dual averaging for convex optimization. In *51st IEEE Conference on Decision and Control*, 2012.
- [34] J. N. Tsitsiklis, D. P. Bertsekas, and M. Athans. Distributed asynchronous deterministic and stochastic gradient optimization algorithms. *IEEE Transactions on Automatic Control*, 31(9):803–812, 1986.
- [35] L. Xiao. Dual averaging methods for regularized stochastic learning and online optimization. *Journal of Machine Learning Research*, 11:2543–2596, 2010.
- [36] L. Xiao and S. Boyd. Fast linear iterations for distributed averaging. *Systems and Control Letters*, 53(1):65–78, Sep. 2004.

- [37] F. Yan, S. Sundaram, S. Vishwanathan, and Y. Qi. Distributed autonomous online learning: Regrets and intrinsic privacy-preserving properties. *IEEE Transactions on Data and Knowledge Engineering*, 25(11):2483–2493, November 2013.